Siglent SDS1000X-E Bandwidth Discussion

Introduction

The Siglent SDS1000X-E series 4-channel models come as 100MHz and 200MHz versions. While the 200MHz model appears particularly tempting, there is still a significant price difference which raises the question whether the higher bandwidth is actually worth the premium.

This article is intended to give some assistance by dealing with the following topics:

1. What are the benefits of the higher bandwidth?
2. Can the SDS1000X-E models cope with 200+ MHz?
3. What are the differences in the 100MHz and 200MHz models frontend?

Siglent were kind enough to offer me an early SDS1104X-E for evaluation and I've also got a SDS1202X-E, which almost certainly has the same frontend as a SDS1204X-E, even though the input capacitance rating is slightly different. This applies to one single channel group, whereas the SDS1104X-E has two of them.

The following chapters describe a number of experiments that should help to shed some light.

Why we need Bandwidth

This question is much harder to answer nowadays, when even entry level scopes like the Siglent SDS1000X-E series start at no lower than 100MHz. The 100MHz bandwidth covers many areas, including audio, RF up to the low VHF band, all sorts of standard logic and 8/16 bit microcontrollers, as well as most peripherals even for the high speed 32 bit MCUs. Who actually needs more?

Well, the ones who really do will of course know why. The more important question would be what are the benefits of 200MHz over 100MHz? First of all, it is not really just 100MHz versus 200MHz. See the tables below for SDS1104X-E vs. SDS1202X-E bandwidth measurements:

	BW [MHz]		
Vert. Gain	1 dB	3 dB	6 dB
0.5mV/Div	74,00	110,00	141,00
1mV/Div	74,00	110,00	141,00
2mV/Div	74,00	110,00	141,00
5mV/Div	74,00	110,00	141,00
10mV/Div	74,00	109,00	141,00
20mV/Div	74,00	111,00	141,00
50mV/Div	74,00	109,00	141,00
100mV/Div	74,00	109,00	141,00
200mV/Div	78,00	109,00	139,00
500mV/Div	76,00	109,00	138,00
1000mV/Div	78,00	109,00	138,00

SDS1104X-E Bandwidth Table

	BW [MHz]		
Vert. Gain	1 dB	3 dB	6 dB
0.5mV/Div	204,00	240,00	290,00
1mV/Div	204,00	240,00	290,00
2mV/Div	203,00	240,00	289,00
5mV/Div	200,00	236,00	285,00
10mV/Div	199,00	236,00	285,00
20mV/Div	201,00	237,00	285,00
50mV/Div	198,00	233,00	283,00
100mV/Div	192,00	227,00	280,00
200mV/Div	200,00	243,00	314,00
500mV/Div	196,00	238,00	310,00

SDS1202X-E Bandwidth Table

It can bee seen, that the differences are bigger than that.

- If you're after accuracy, then the -1dB bandwidth is roughly 74MHz vs. 200MHz, factor 2.7
- The usual benchmark -3dB and -6dB bandwidths are much closer to factor 2 though

The actual -1dB bandwidth is pretty much identical with the specified bandwidth in the 200MHz model, but only some 75% of that for the 100MHz model.

While we obviously need more bandwidth when dealing with higher frequencies, we might also want it even with low frequency signals if they happen to include high frequency content (e.g. fast transitions) and we're interested in a faithful rendition of the original waveform. A spectrum analyzer can be used in order to learn about the bandwidth needed to get a high fidelity representation of the signal on the screen. Look at the 2MHz square wave from my 60MHz AWG as an example:

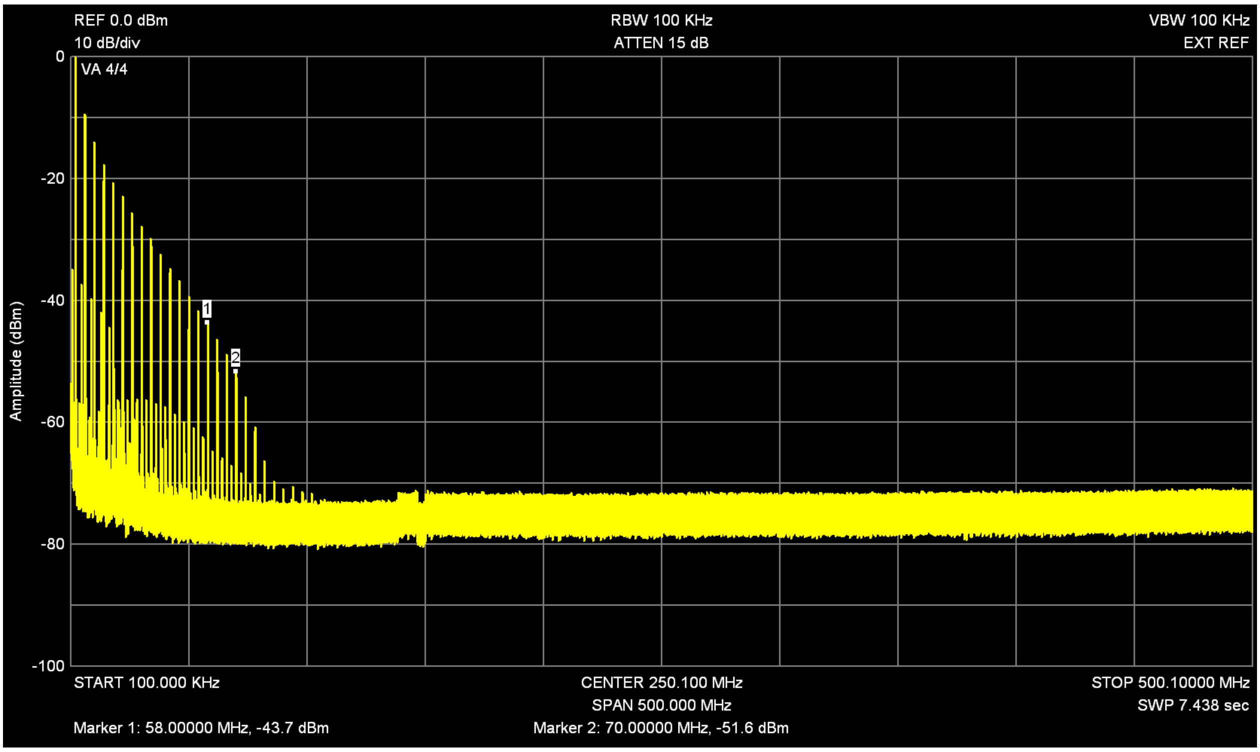

UTG2062A_Square_2MHz_Spectrum

When looking at this spectrum, we know at first glance that 70MHz bandwidth would be plenty to give us an almost perfect representation of the real signal.

It is a little different for a 2MHz signal from my ancient 100MHz pulse generator, which has some serious VHF-power transistors in its output stage:

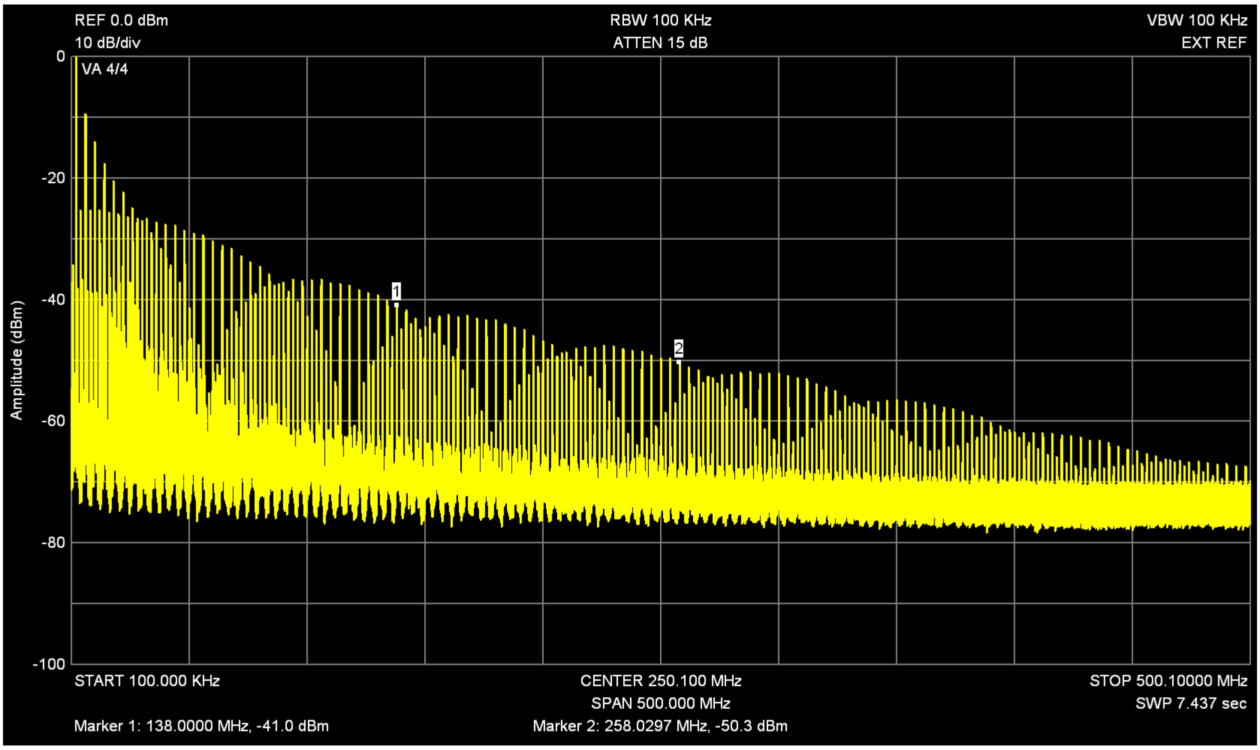

PM5771_Square_2MHz_Spectrum

Here we'd probably need 200MHz bandwidth to keep the total error below 1%.

Some practical examples will be shown later in the "Square Wave Test" section.

As a verdict for this chapter it appears to be the same old story again: the more the merrier. Yet this can also have its drawbacks. To utilize higher bandwidth, we need better probes alongside with better probing techniques, better cables and connectors – which all sums up in additional cost and effort. Thankfully, this is not yet an issue if we're talking about just 200MHz. But there might be other pitfalls, see next chapter…

Why we need Sample Speed

High bandwidth is great, but without high sample speed it could create more headaches than benefits. At all times, there have been digital scopes with a bandwidth higher than the maximum sample speed. These scopes use equivalent time sampling in order to digitize strictly periodic signals up to the full input bandwidth. Today, this technique is mostly limited to multi GHz digitizers and not often found in modern general purpose scopes anymore. These are true real time (as opposed to equivalent time) scopes now, which appears to be great from a user's perspective, as there's no need to worry about the limits of the real-time bandwidth. But now we might start worrying about aliasing…

Everyone knows the sample theorem. In short, we can faithfully reconstruct a strictly bandwidth limited signal as long as we sample it at a rate that is at least twice the highest input frequency. In other words, the highest input frequency must not exceed half the sample rate, which is also known as the Nyquist frequency. Does that mean we need only 400MSa/s sample rate for a 200MHz scope?

Of course not. The sample theorem is just theory assuming a strictly bandwidth limited signal, digitized by a perfectly linear ADC with infinite resolution (in fact, the sample theorem only deals with quantization of time, but not magnitude) and finally the reconstruction with a perfect brick-wall filter.

It should be clear that none of these conditions can be met in any practical device. We can get a reasonably bandwidth limited signal out of a low distortion sine wave generator, but every other signal source might well have significant harmonics and spurs above the Nyquist frequency. We'd need an input filter with ideal brick-wall characteristic to remove that "forbidden" signal content, this is also known as anti-aliasing filter. Yet I very much doubt that any real DSO actually has one. Of course they do not, because the bandwidth limit of a scope traditionally is supposed to have some sort of Gaussian response and it's actually the natural roll-off of the input amplifier that ultimately determines the input bandwidth. No one wants a brick-wall filter at the scope input for various reasons and if manufacturers actually fit some filter (in high end scopes), then this is for equalization of phase response and group delay, which is quite the opposite from what a brick-wall filter does. In fact, the only way to fulfill this very first requirement of keeping out signals above Nyquist is oversampling to an extent that the natural roll-off of the input amplifiers yields sufficient attenuation at and above the Nyquist frequency.

The next problem is the quantizing of the signal magnitude. The fact that we only have 256 discrete steps in an 8-bit ADC does not invalidate the sample theorem, but every ADC has some DNL and INL (differential and integral nonlinearity). So the ADC itself produces harmonics and unwanted mixer products, some of them above the Nyquist frequency – and these signal components cannot be filtered out! But then, a good ADC with < 1 LSB INL shouldn't cause any visible problems.

Finally we need reconstruction. After all, we want to create a contiguous signal out of the discrete samples again. Linear interpolation was an early method, still offered by most scopes today. It requires little processing power, but does a lousy job, as it cannot be called a reconstruction filter at all. Consequently, signal rendition will only be reasonable as long as the sample rate is at least ten times the signal frequency. Sin(x)/x aka Sinc is a true reconstruction function that would act as a brick-wall filter – if it could process an infinite number of data points. Quite obviously this is not going to happen in any scope, let alone at fast timebase settings, where the actual visible record consists just of a handful of samples. Yet the results can be surprisingly good and a Sinc filter that uses all available samples manages to reasonably reconstruct a signal at 40% of the sample rate, which is equivalent to 80% of the Nyquist frequency.

As a result, it is actually possible to reconstruct a 200MHz signal that has been sampled at only 500MSa/s. But any higher than that will inevitably lead to distortion and modulation effects. Furthermore, a 200MHz frontend will not stop frequencies above 250MHz i.e. the Nyquist frequency at 500MSa/s from reaching the ADC, thus causing aliasing artifacts.

Now let's check how all these requirements are met in the 200MHz SDS1202X-E and 100MHz SDS1104X-E when both channels in a group are enabled and the sample rate is limited to 500MSa/s.

Aliasing in the SDS1202X-E

After the explanations given earlier in this article, we cannot expect any wonders. This 200MHz scope has an actual 3dB bandwidth of some 240MHz and since the -6dB bandwidth is at least 280MHz, we can safely state that there is barely any attenuation at the 250MHz Nyquist frequency for 500Msa/s. Of course, this should not be a big problem with only one channel enabled, because we have 1GSa/s sample speed then and Nyquist frequency appears sufficiently high at 500MHz.

To estimate the attenuation at higher frequencies, we start at 10MHz in order to establish a reference value. 1GSa/s is used for these tests, because we want to see the attenuation and need the scope working at frequencies higher than 200MHz.

Sine_10MHz_BW237MHz_1GSa

Now let’s have a look at a signal at the nominal bandwidth limit of 200MHz.

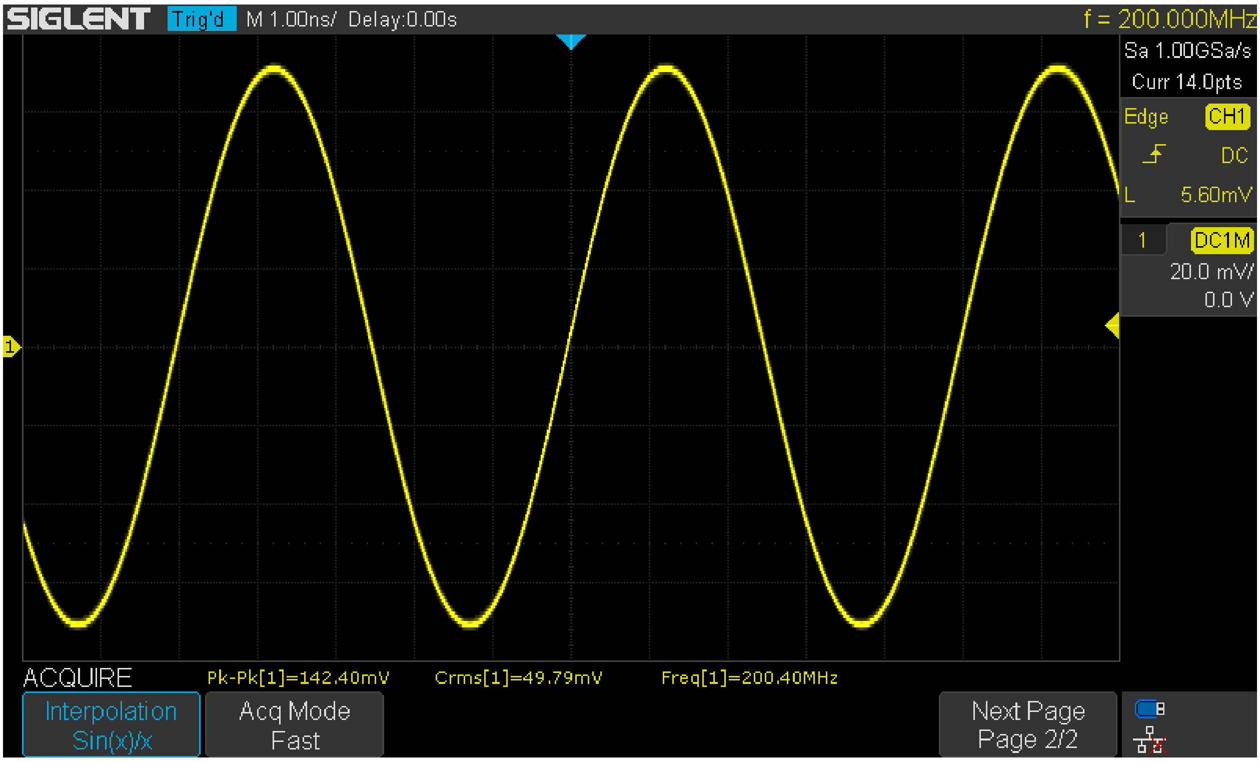

Sine_200MHz_BW237MHz_1GSa_SinX

We can barely see any attenuation. Even with 1GSa/s, we would get just garbage if we had no sin(x)/x reconstruction available. See the following screenshot with just linear interpolation at that frequency:

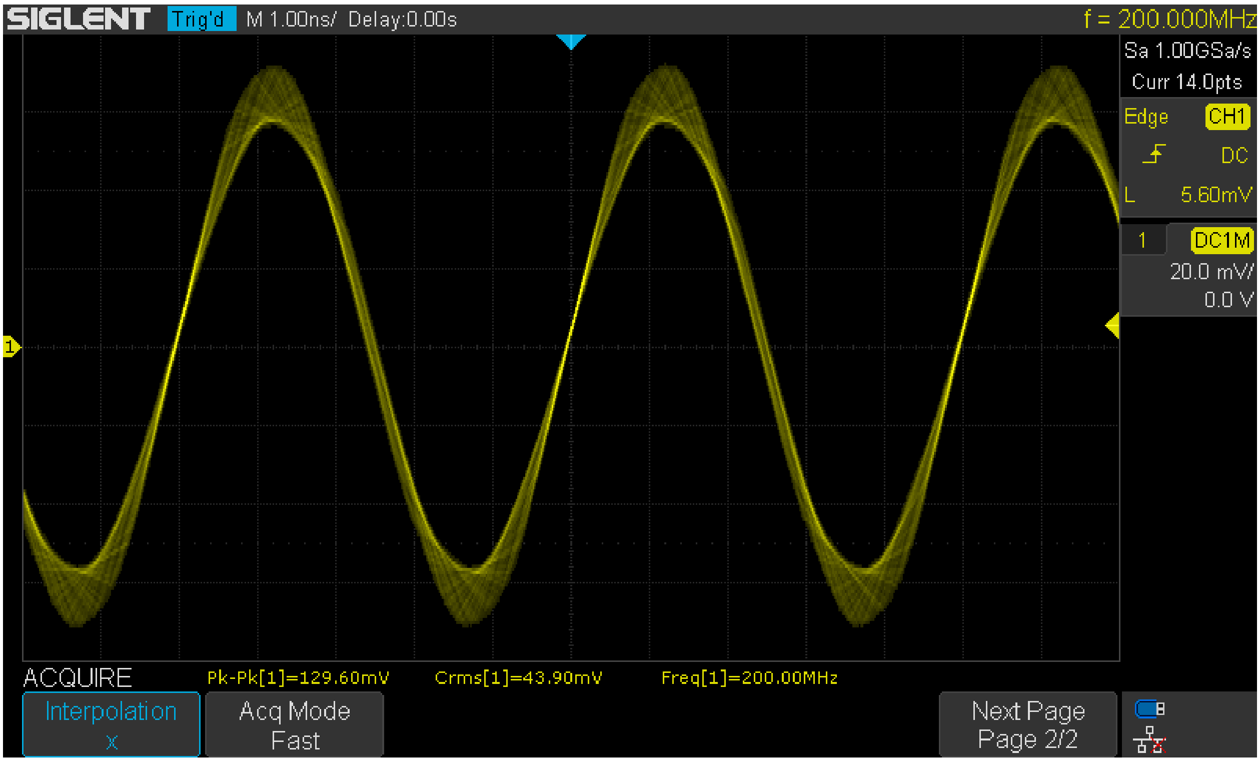

Sine_200MHz_BW237MHz_1GSa_X

Now let's check the signal at 250MHz, which is the Nyquist frequency for 500MSa/s.

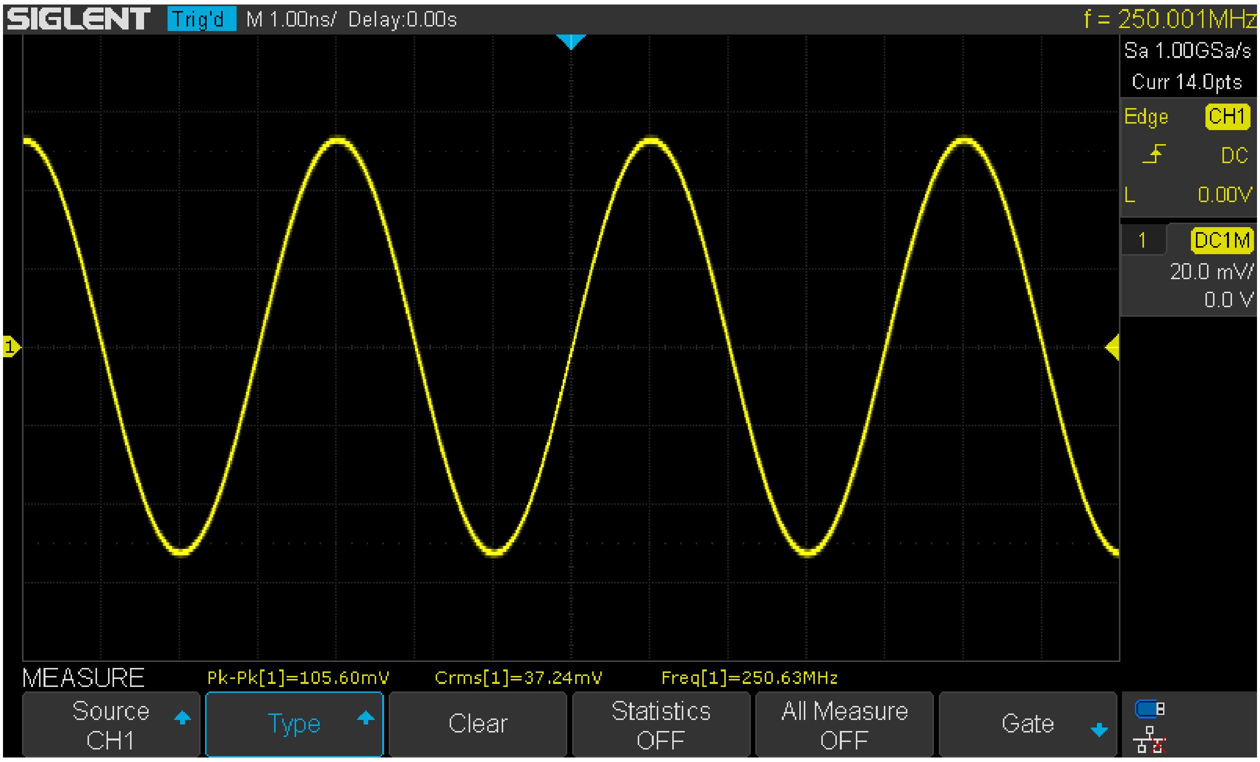

Sine_250MHz_BW237MHz_1GSa

There is still not much attenuation; 105/160 = 0.656 = -3.66dB. Quite obviously, there is absolutely no protection against aliasing. Now let's check the signal at 500MHz – the Nyquist frequency for 1GSa/s.

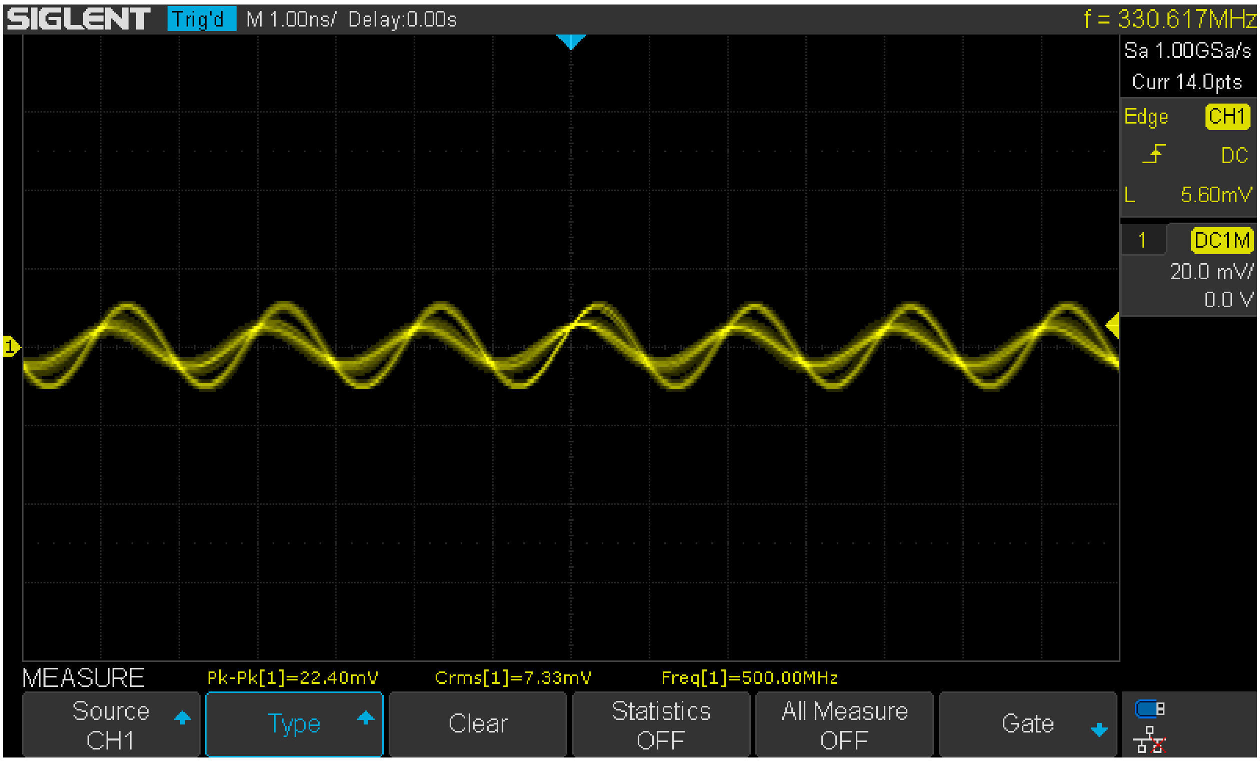

Sine_500MHz_BW237MHz_1GSa

Unsurprisingly, we get a garbage trace at the Nyquist frequency, but we can still see the maximum amplitude of the modulated waveform – and this is unexpectedly high at about 22mV. Attenuation is weak: 22.4/160 = 0.14 = -17dB. This is barely enough to prevent signal distortion if we have an input signal with a significant amount of energy above 500MHz. Thankfully, this is not very likely to happen in an average electronic project where entry level scopes are in use – and the folks dealing with really fast signals will certainly know that an entry level scope is not going to be up to the task.

On the other hand, most users on a budget owning this scope would be proud that the advertised bandwidth is more like 1dB rather than 3dB and that signals up to 420MHz can be displayed without major issues. The frequency counter still works and the automatic time measurements are correct, even though they now have to deal with values of 1ns and less.

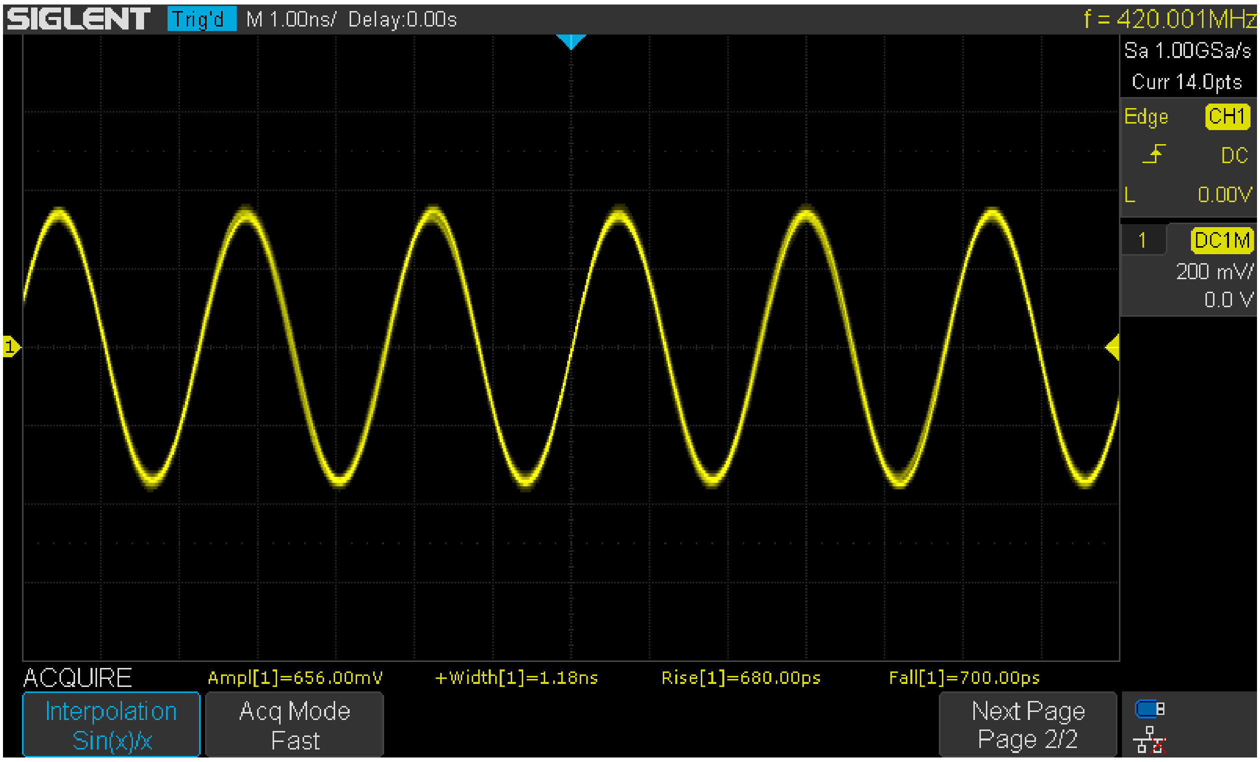

Sine_420MHz_BW237MHz_1GSa

Reconstruction in the SDS1202X-E

This scope did not shine in the (anti) aliasing section, as there is absolutely no attenuation worth mentioning at the Nyquist frequency. Furthermore the frontend appears to have a rather slow roll-off, so that even at 500MHz a significant portion of the signal will pass through. Let's see if at least the reconstruction filter makes the best out of the few data points at high frequencies.

In actual fact, this question has already been answered by the previous test with the 420MHz signal, where only slight signs of amplitude modulation have become visible. But now we want to look at that a bit more systematically.

Let's start at the specified bandwidth of 200MHz and a sample rate of only 500MSa/s, as with all channels enabled. First screenshot is with sin(x)/x reconstruction, 2nd one with linear interpolation. What a difference!

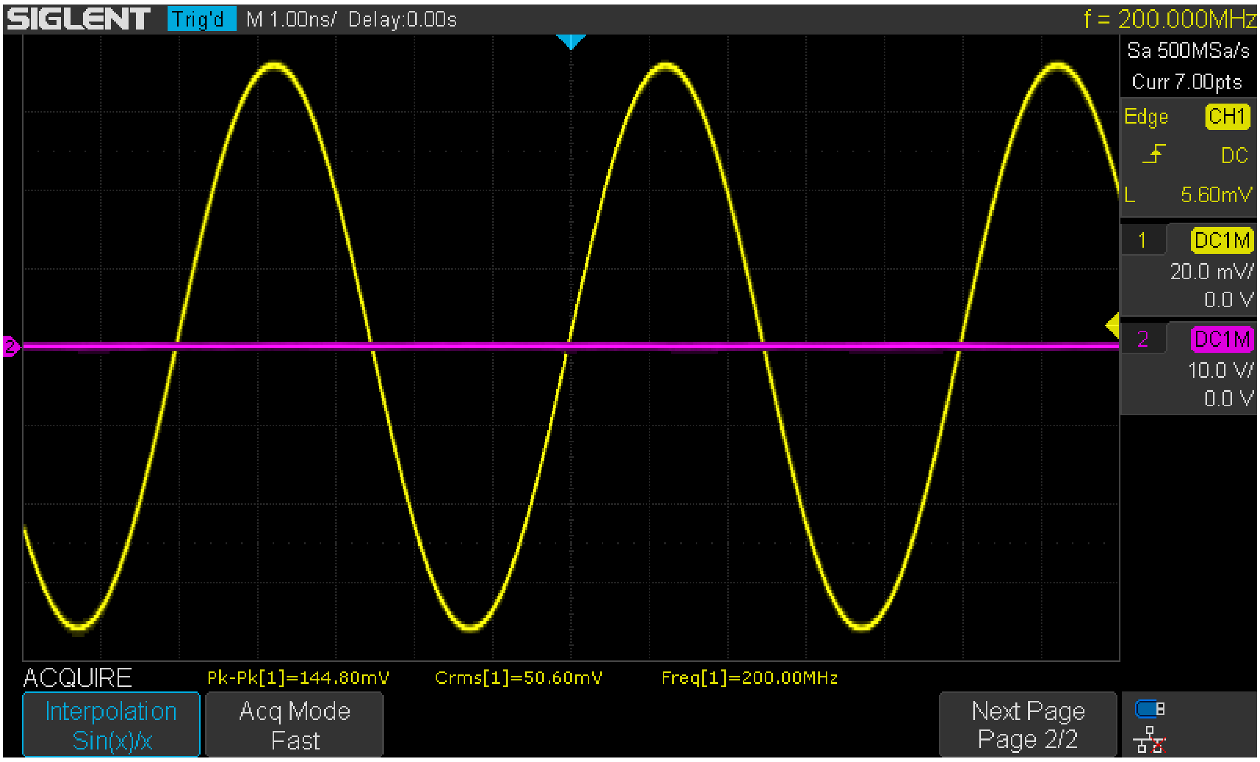

Sine_200MHz_BW237MHz_500MSa_SinX

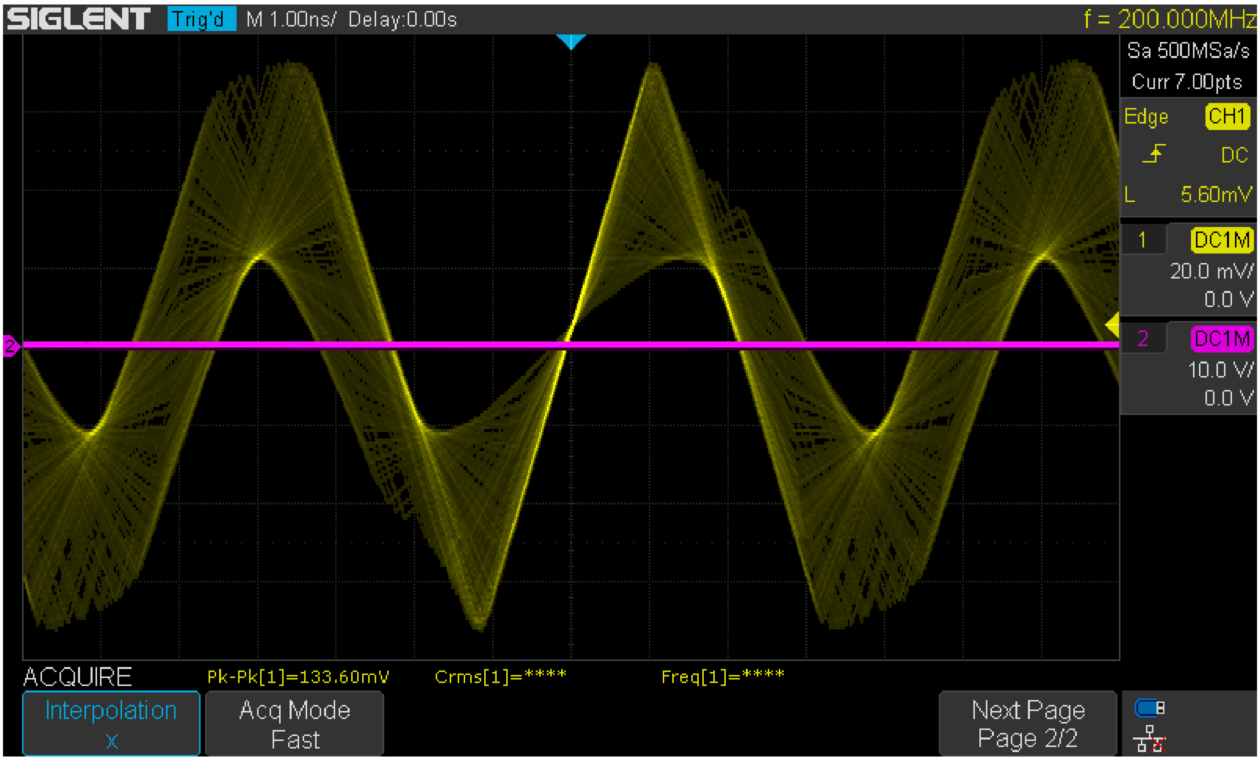

Sine_200MHz_BW237MHz_500MSa_X

Sin(x)/x reconstruction is perfectly fine, whereas linear interpolation gives awful results. But we can always get an artifact-free rendition by using dots display mode. Well, at frequencies that high (and only 7 samples per screen width), we don’t get a contiguous trace anymore, but that’s really a very minor problem.

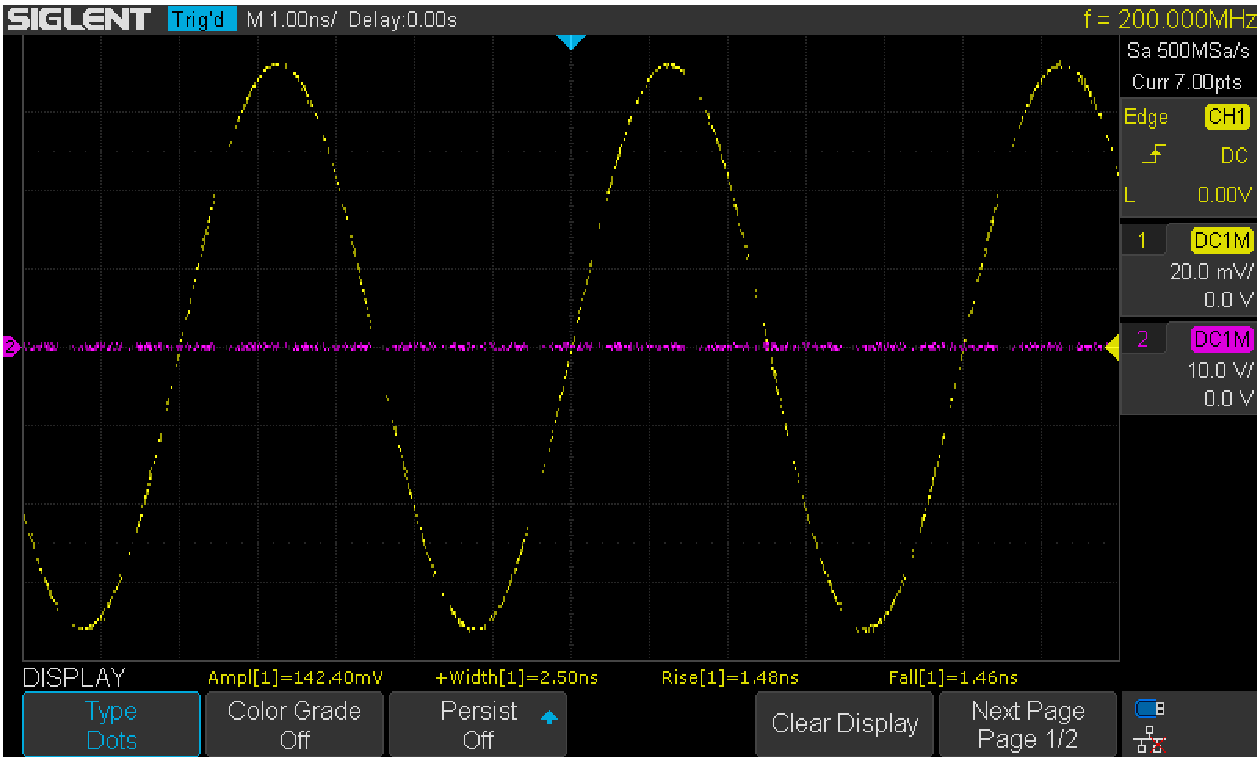

Sine_200MHz_BW237MHz_500MSa_Dots

As we could see, dots mode as well as sin(x)/x reconstruction yielded flawless renditions at a signal frequency 80% of Nyquist. Now we want to see what it looks like if we go any higher.

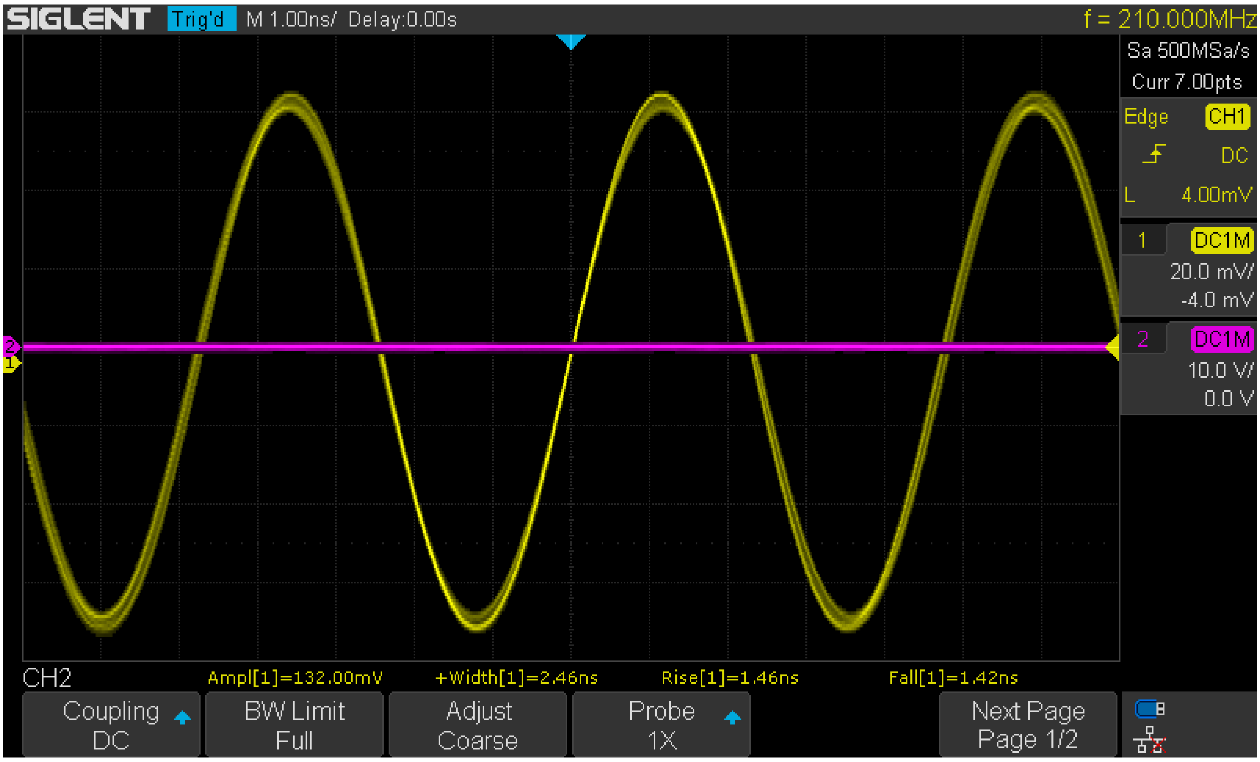

Sine_210MHz_BW237MHz_500MSa_SinX

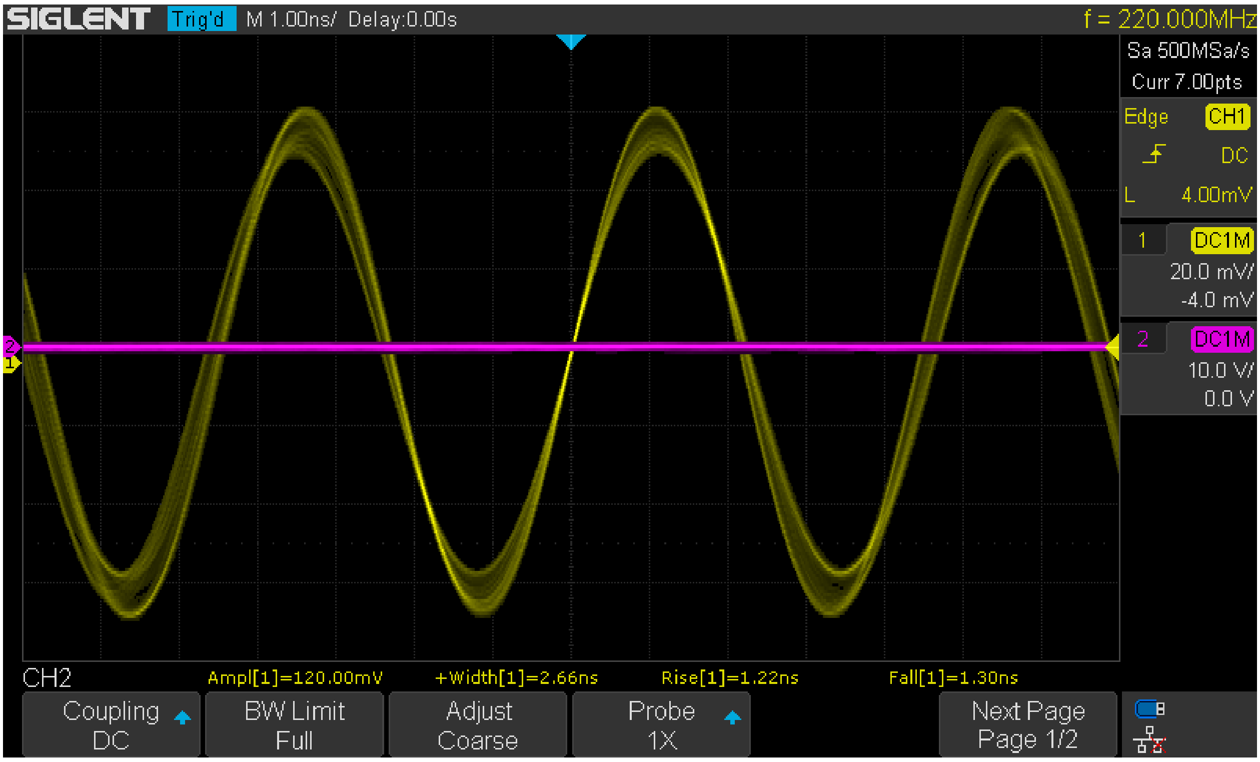

Sine_220MHz_BW237MHz_500MSa_SinX

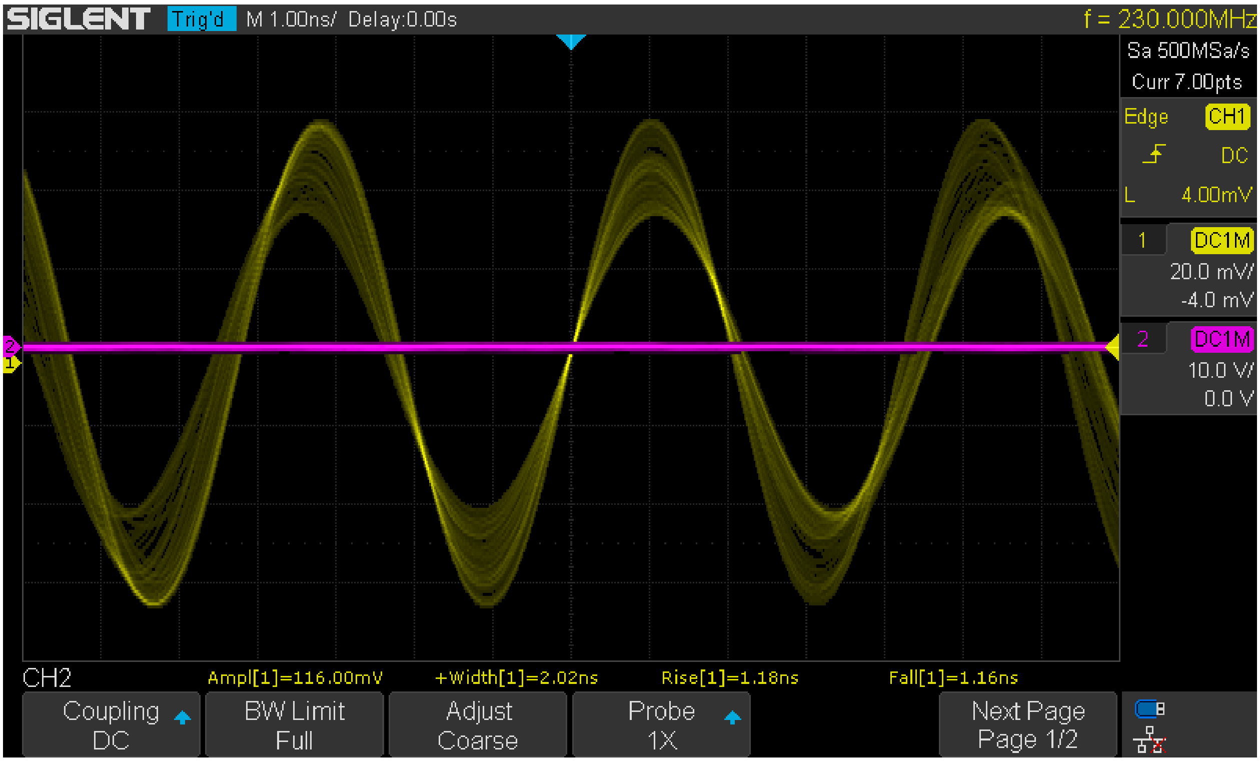

Sine_230MHz_BW237MHz_500MSa_SinX

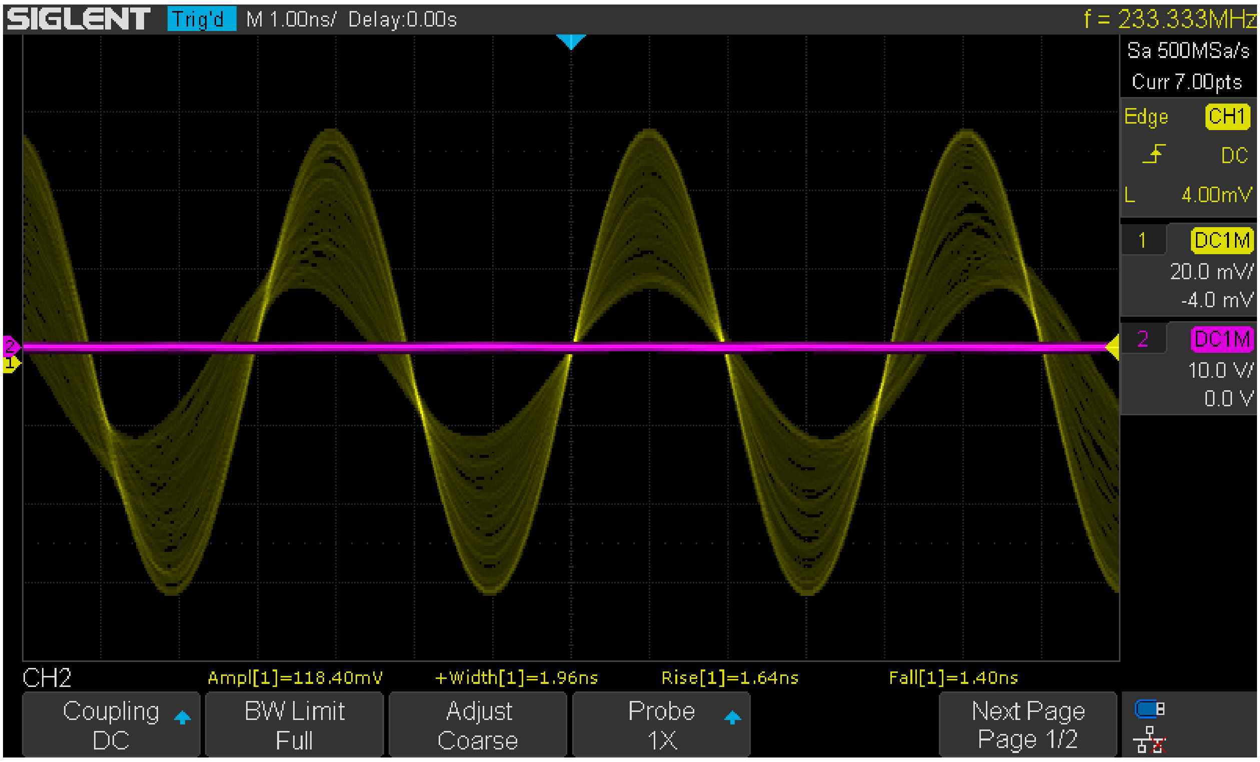

Sine_240MHz_BW237MHz_500MSa_SinX

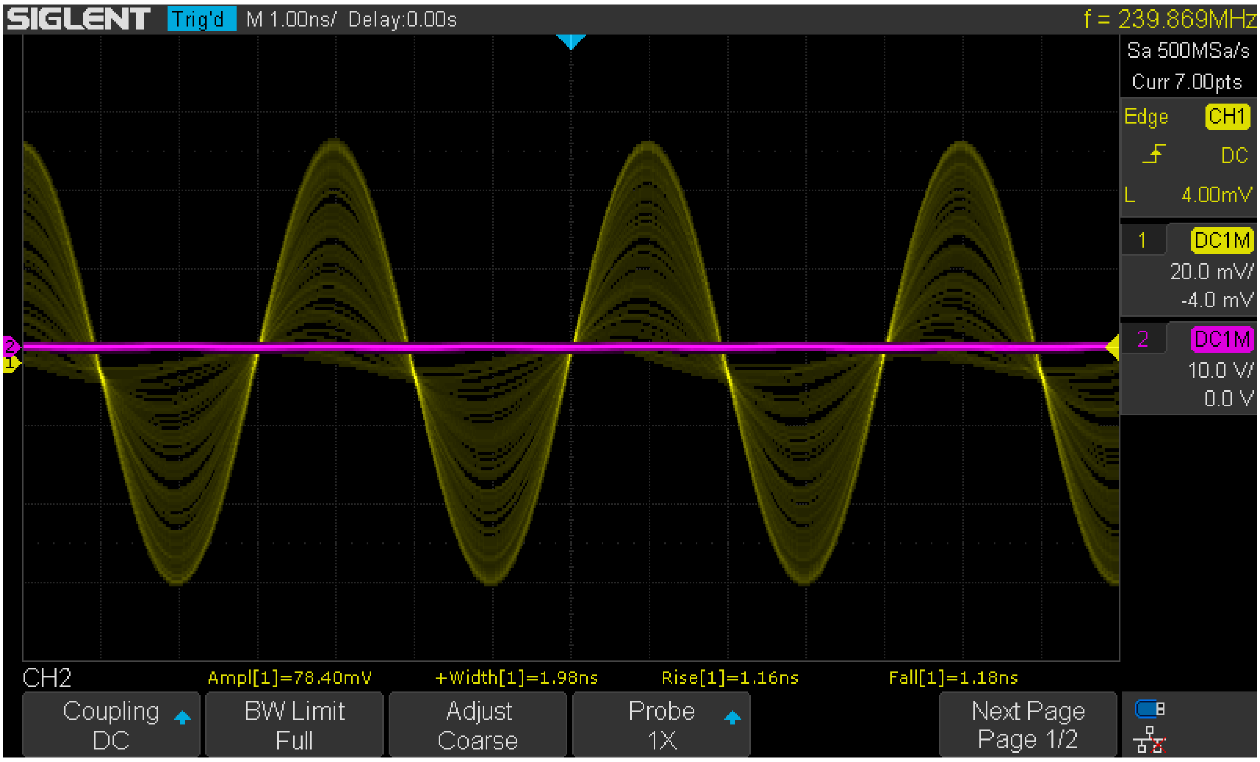

Sine_250MHz_BW237MHz_500MSa_SinX

The screenshots speak for themselves. With only 7 samples for the entire record, sin(x)/x reconstruction inevitably fails at frequencies any higher than 80% of Nyquist. These screenshots also demonstrate the value of high waveform capture rates and intensity grading that actually works.

Aliasing in the SDS1104X-E

Aliasing shouldn't be a problem at all for the 100MHz scope. The 6dB bandwidth is some 140MHz and we expect a 3rd order roll-off above that point, so we should get some 23dB attenuation at the Nyquist frequency for 500MSa/s.

To estimate the attenuation at higher frequencies, we start at 10MHz in order to establish a reference value. 1GSa/s is used for these tests, because we want to see the attenuation and need the scope working at frequencies higher than 200MHz.

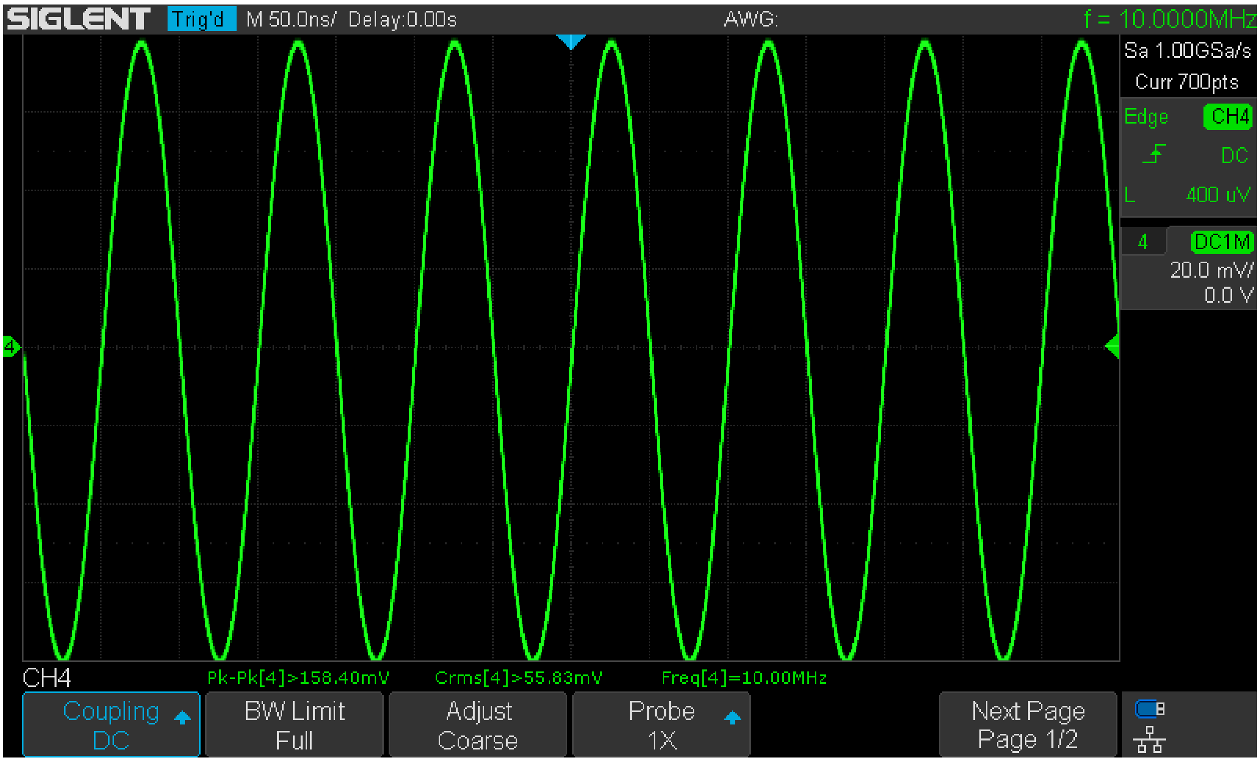

Sine_10MHz_BW111MHz_1GSa

Now let's have a look at a signal at the nominal bandwidth limit of 100MHz.

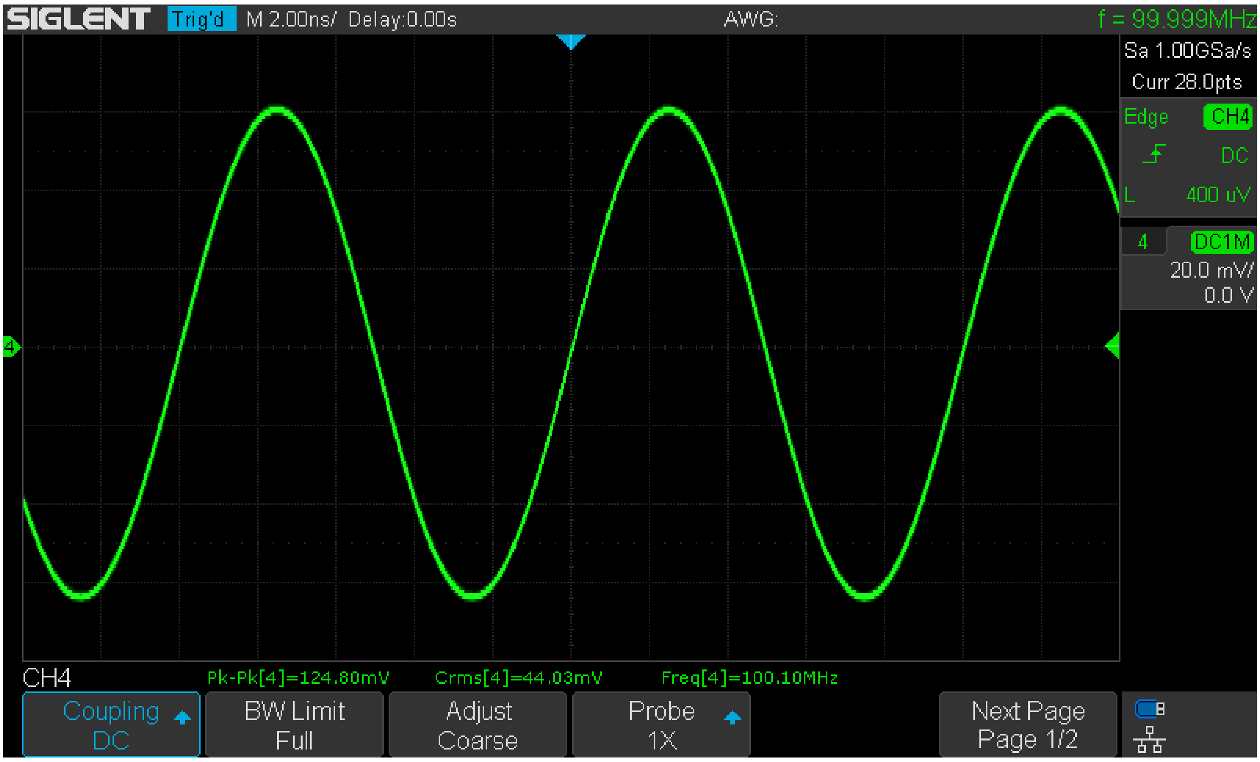

Sine_100MHz_BW111MHz_1GSa

We don't get much attenuation yet, but it's still more than the SDS1202X-E at 200MHz. So let's check 200MHz then:

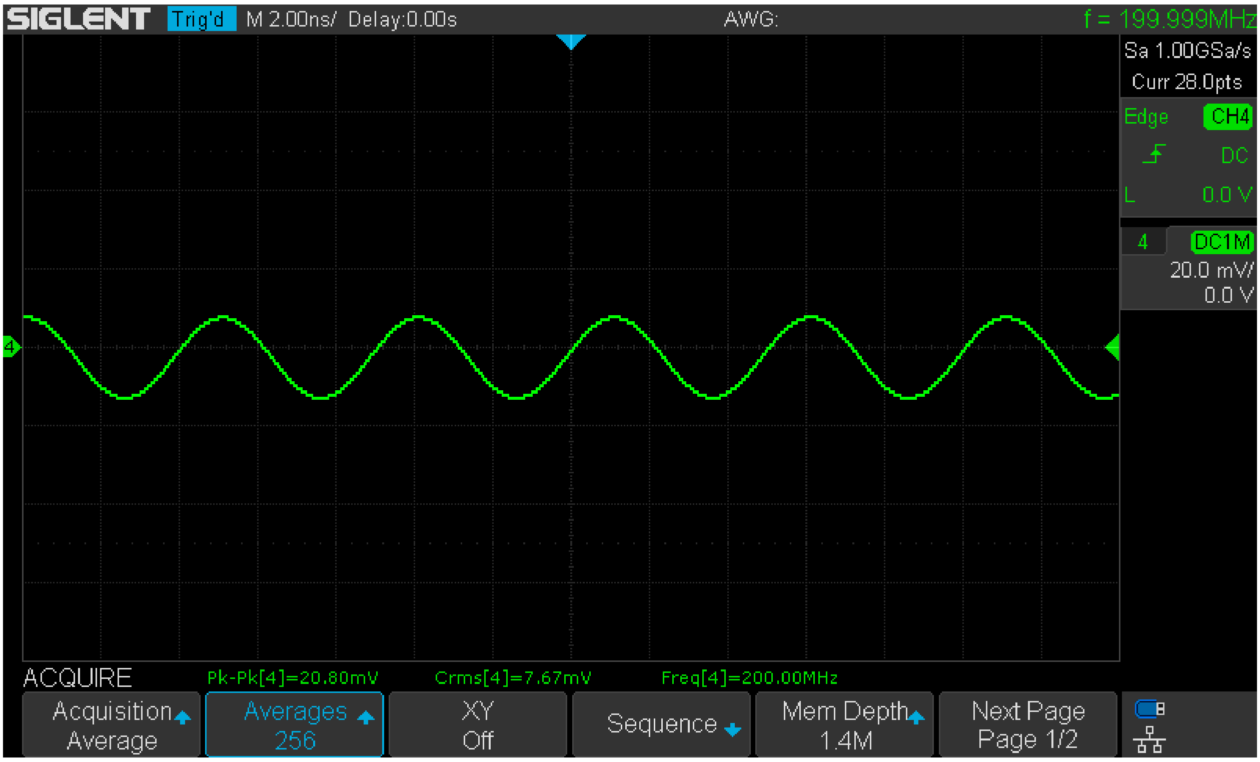

Sine_200MHz_BW111MHz_1GSa

There is already some noticeable attenuation: 20/158 = 0.127 = -18dB
Now let's check the signal at the 250MHz Nyquist frequency for 500MSa/s.

Sine_250MHz_BW111MHz_1GSa

The signal vanishes in the noise floor. From the automatic measurements, the attenuation has to be at least 3.2/158 = 0.02 = -34dB. While this would certainly prevent any visible signal distortion, it is way more than expected. An additional attenuation of 16dB for the frequency step from 200MHz to 250MHz is not plausible for an ordinary low-pass filter unless it had Cauer characteristics. But then again, we certainly wouldn't expect a Cauer filter in an oscilloscope frontend! We will further investigate this in the last chapter - *The 100MHz Puzzle*.

Reconstruction in the SDS1104X-E

There's nothing to show here, since reconstruction works exactly the same as in the SDS1202X-E.

Square Wave Test

For a real world comparison, a square wave with fast transients <100ps has been used. A worst case scenario has been set up by using a fundamental frequency of 50MHz, so that the 5th harmonic is sitting exactly at the Nyquist frequency for 500MSa/s. Because of the fast transients there is a lot of energy above the Nyquist frequency. That's some really nasty spectrum, even though it's just a square wave and not a pulse train:

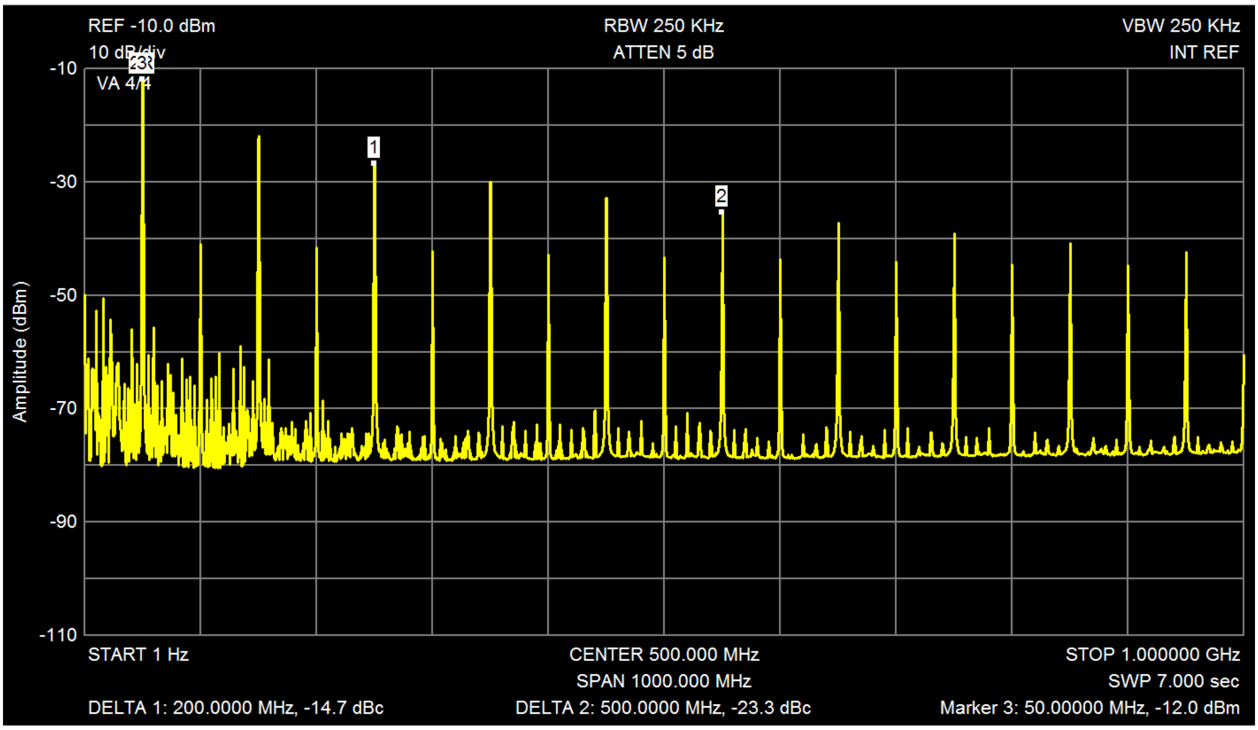

Square_50MHz_Spectrum

First let's have a look at the waveform with a scope that provides 350MHz bandwidth – of course, this is not nearly enough and a (at least) 1GHz scope would be in order for this, but it still shows a noticeable difference to the 237MHz SDS1202X-E.

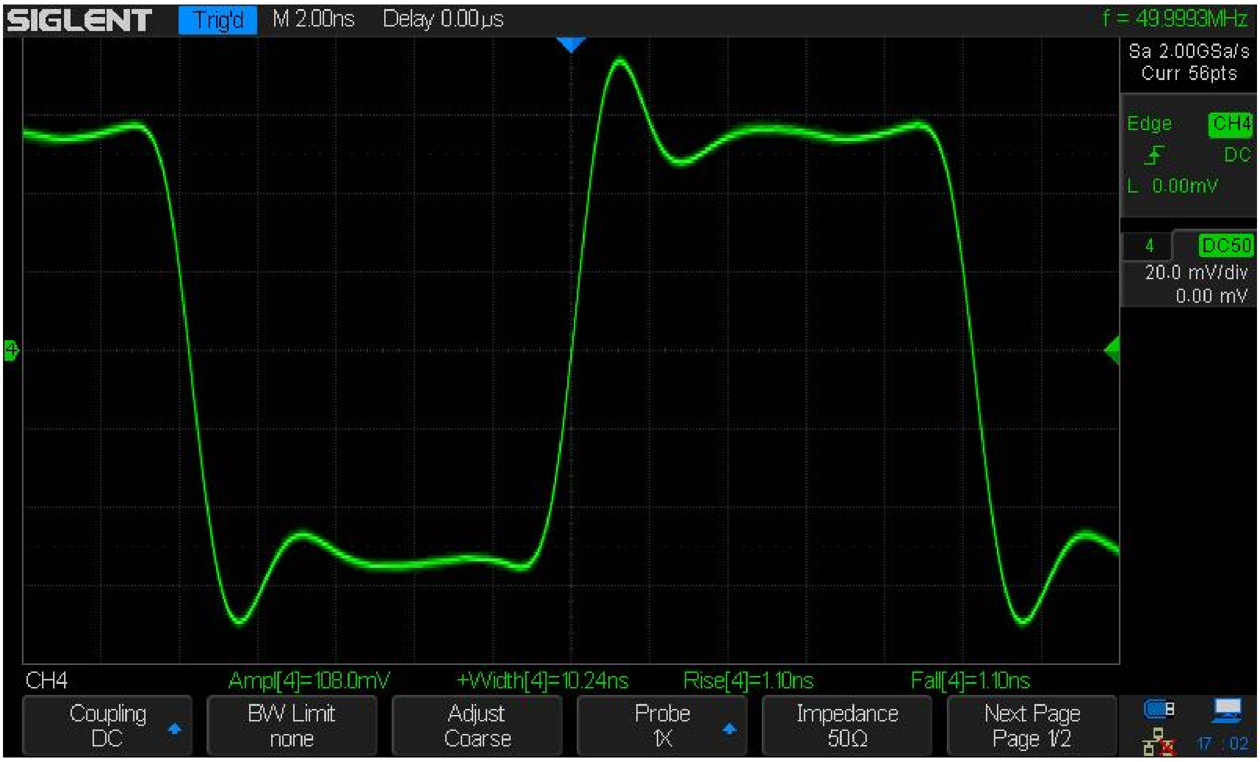

Square_50MHz_BW351MHz_2GSa

Let's keep the measurements in mind and compare them with the subsequent screenshots. Now the SDS1202X-E:

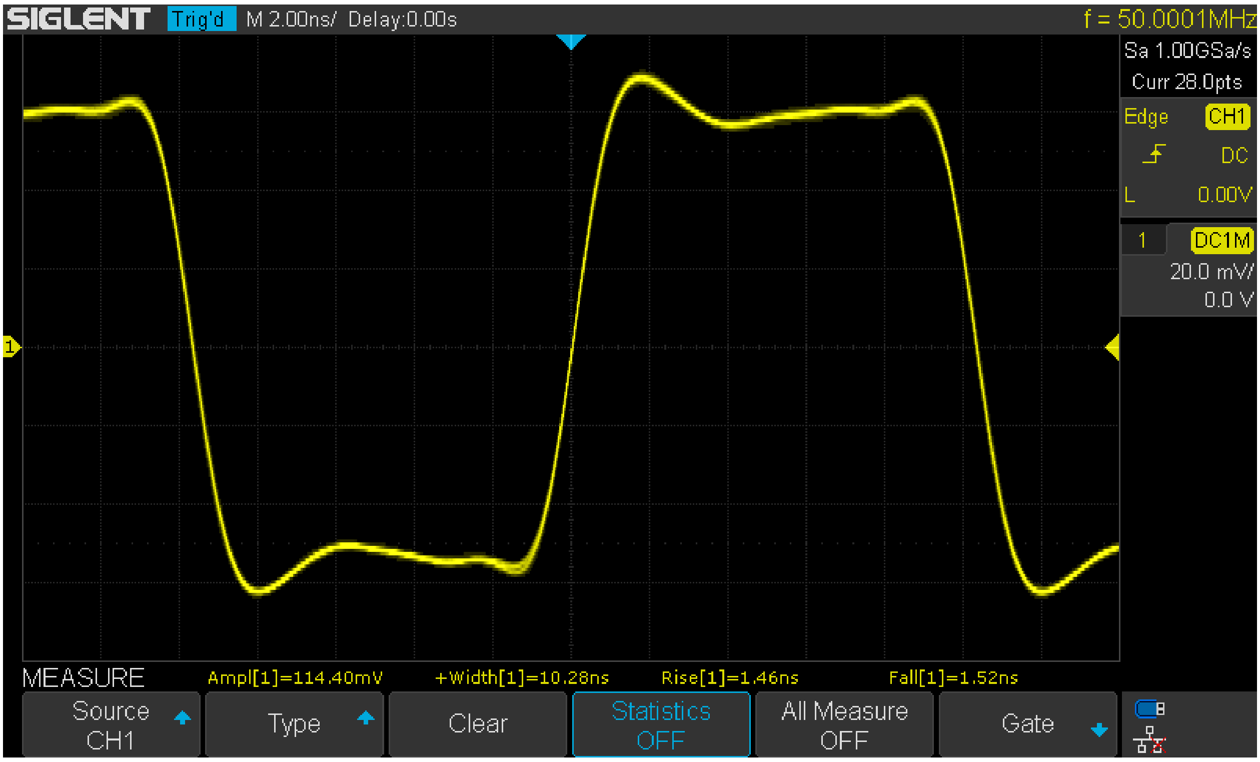

Square_50MHz_BW237MHz_1GSa

237MHz vs. 351MHz – that’s a very noticeable difference already and the trace gets a little fuzzy at some spots even at 1GSa/s, yet we get an idea what the signal might actually look like. Now for the SDS1104X-E:

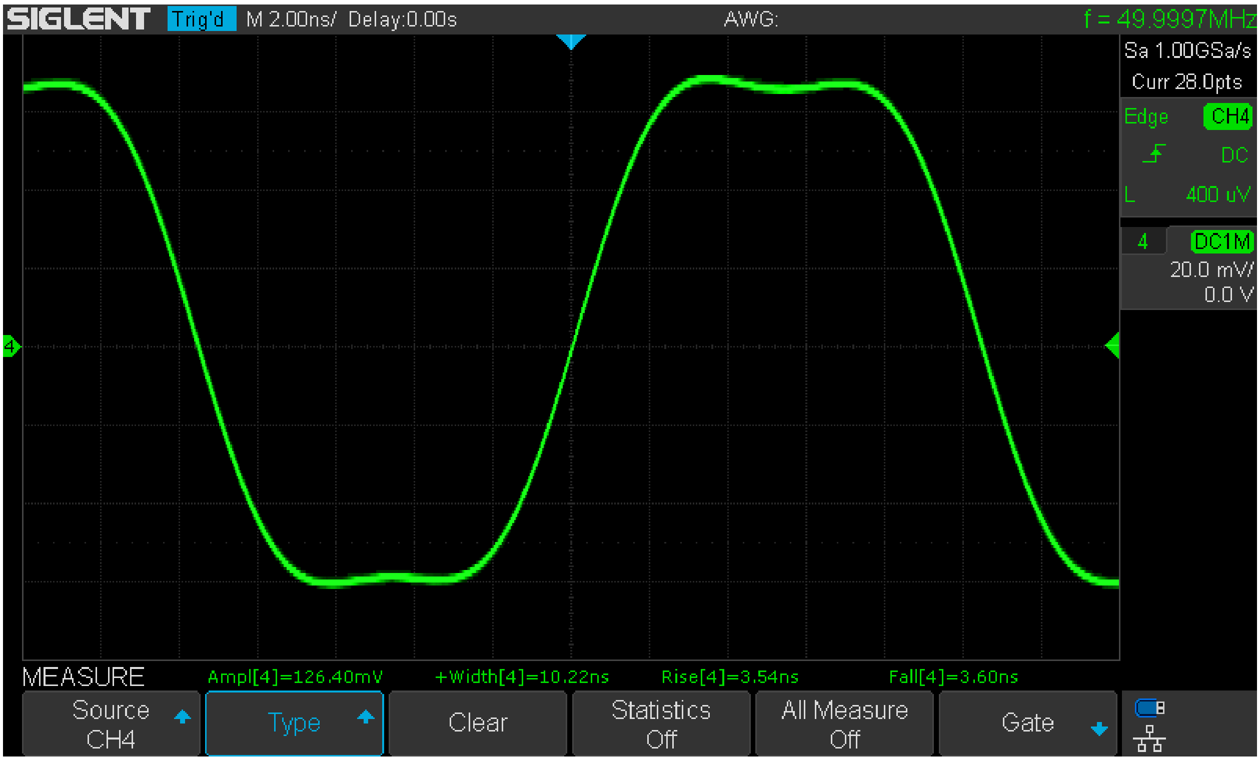

Square_50MHz_BW111MHz_1GSa

All the signal detail is gone now and the trace looks rather clean. However, this test confirms that we can never have enough bandwidth in some situations. Yet the pulse width measurements are pretty much identical, just the transitions differ as they naturally just show the scope’s own rise time.

With a fast signal like this, first signs of tiny artifacts have already shown up in the SDS1202X-E even at 1GSa/s. So what will it look like at half that sample rate? First with sin(x)/x reconstruction:

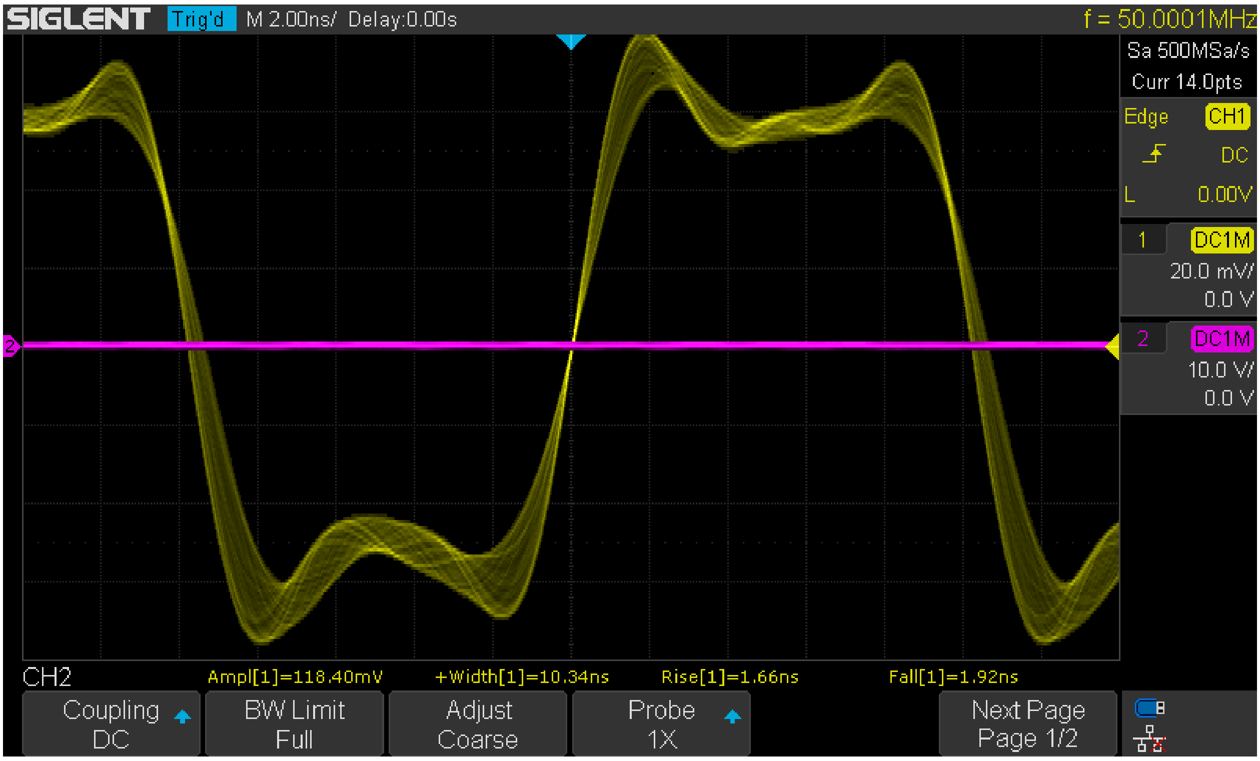

Square_50MHz_BW237MHz_500MSa

Oh yes, these are the nasty effects of aliasing. Now let's have a look at dots display mode:

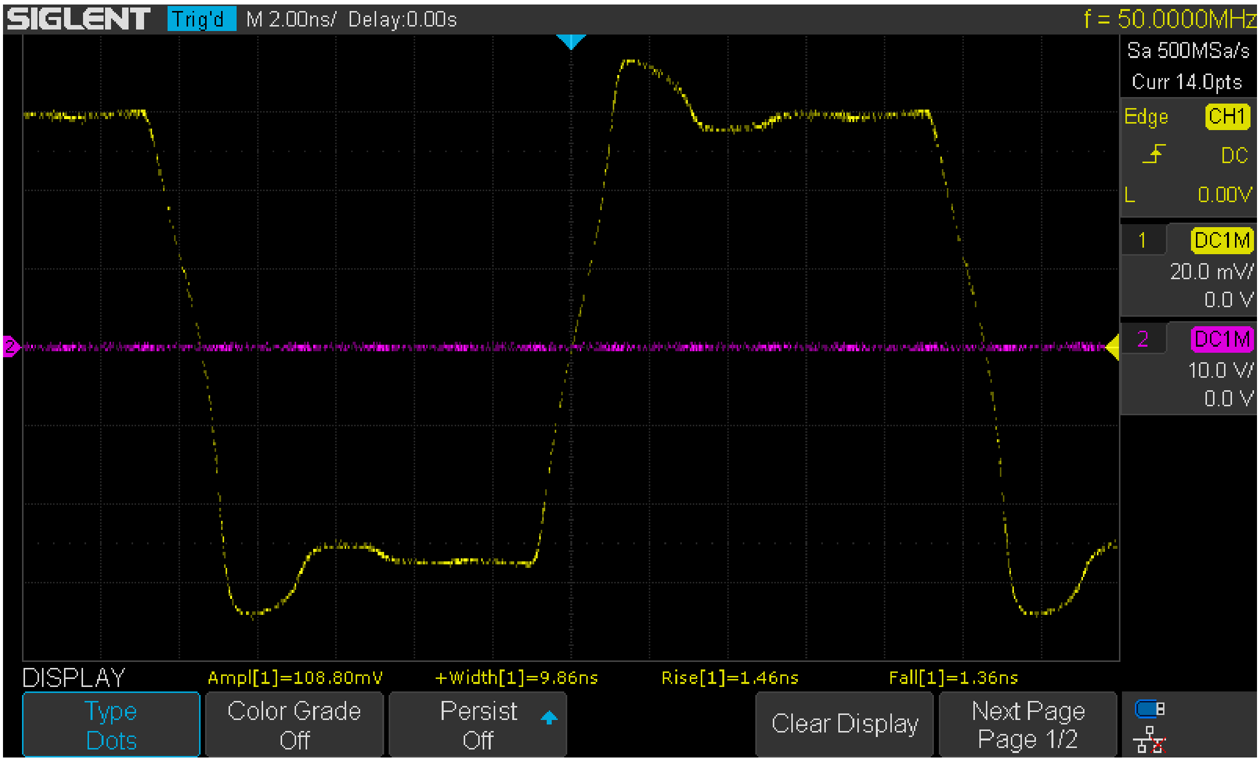

Square_50MHz_BW237MHz_500MSa_Dots

That looks ugly as well and the kinky waveform clearly hints on the aliasing artifacts we're seeing. Yet there is now absolutely no jitter and the automatic measurements generally show lower values because there are no reconstruction errors any more.

So what are the consequences?

As could be demonstrated, even with a very fast signal and vectors display mode, the SDS1202X-E works reasonably well with sin(x)/x reconstruction at 1Gsa/s, which in turn means that only one channel can be used. That's very limiting indeed and the SDS1204X-E would make more sense here, because it gives us two fully usable channels in situations like this.

On the other hand, when dealing with super fast signals like this we might not even be interested in the true waveform – which we're not going to see anyway, even with a 500MHz scope – and then the 200MHz SDS1000X-E series actually appears to have the major drawback of delivering ugly pictures with lots of edge jitter like we've seen before. Then again, we can always use dots mode as a last resort, which won't change the ugly waveform rendition but eliminate the AM and jitter problems. Always keep in mind it's only a problem with very fast signals and even then, automatic time measurements aren't far off either way.

Now what about the bandwidth limited 100MHz version of this scope? As we've seen earlier, attenuation at and above Nyquist appears very high so we would not expect any aliasing problems whatsoever when using a SDS1104X-E at timebase settings up to 1ms/div, where a sample rate of 500MSa/s can be maintained even with all four channels in use. Now let's try that:

Square_50MHz_BW111MHz_500MSa

Oops – this doesn't look as bad as the 200MHz scope, but is still not what we'd expected. Why don't we get a clean trace? This calls for further investigation.

The 100MHz Puzzle

We have investigated the behavior of the Siglent SDS1000X-E series, both the 100MHz and 200MHz versions and the very last test of all things yielded quite unexpected results.

The tests should have proven that:

- Frequency response is down at least 34dB at the Nyquist frequency of 250MHz for the SDS1104X-E
- The SDS1000X-E do an excellent job at sin(x)/x signal reconstruction up to 40% of the Nyquist frequency and Dots mode can always be used as a last resort to get rid of any reconstruction issues

Given all that, it would just not be possible to get a fuzzy trace rendition like in the last test, even with the fastest possible signal. Consequently, one of the two assumptions listed above must be wrong.

Let's start with the reconstruction – this just cannot be wrong as the whole process is very transparent. We can see the dots and we can also see the result of the sin(x)/x reconstruction with these very dots. There really shouldn't be any possible way to fool us. So let's review the SDS1104X-E with the fast transients in dots display mode again:

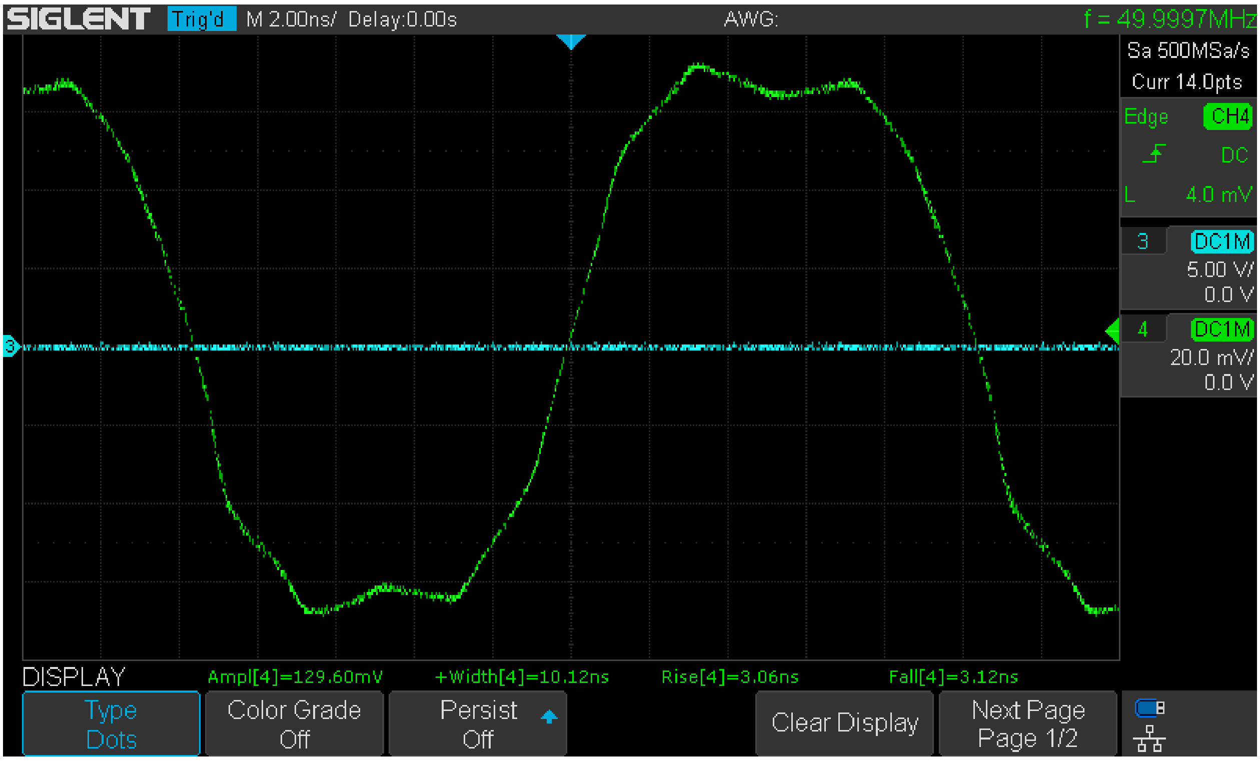

Square_50MHz_BW111MHz_500MSa_Dots

Well, that's the well know aliasing artifacts, not as bad yet similar to the SDS1202X-E. This is a strong hint that the assumption of a high attenuation at the Nyquist frequency must be wrong.

How can we prove that?

Well, let's turn an SDS1202X-E into an SDS1102X-E by just adding a 100MHz filter to the input. Since the measurements suggest an approximate 3rd order filter slope between 100 and 200MHz, I've quickly tinkered a 3rd order 50Ω filter fitted into a small metal box with BNC connectors for this purpose. The frequency response looks like this:

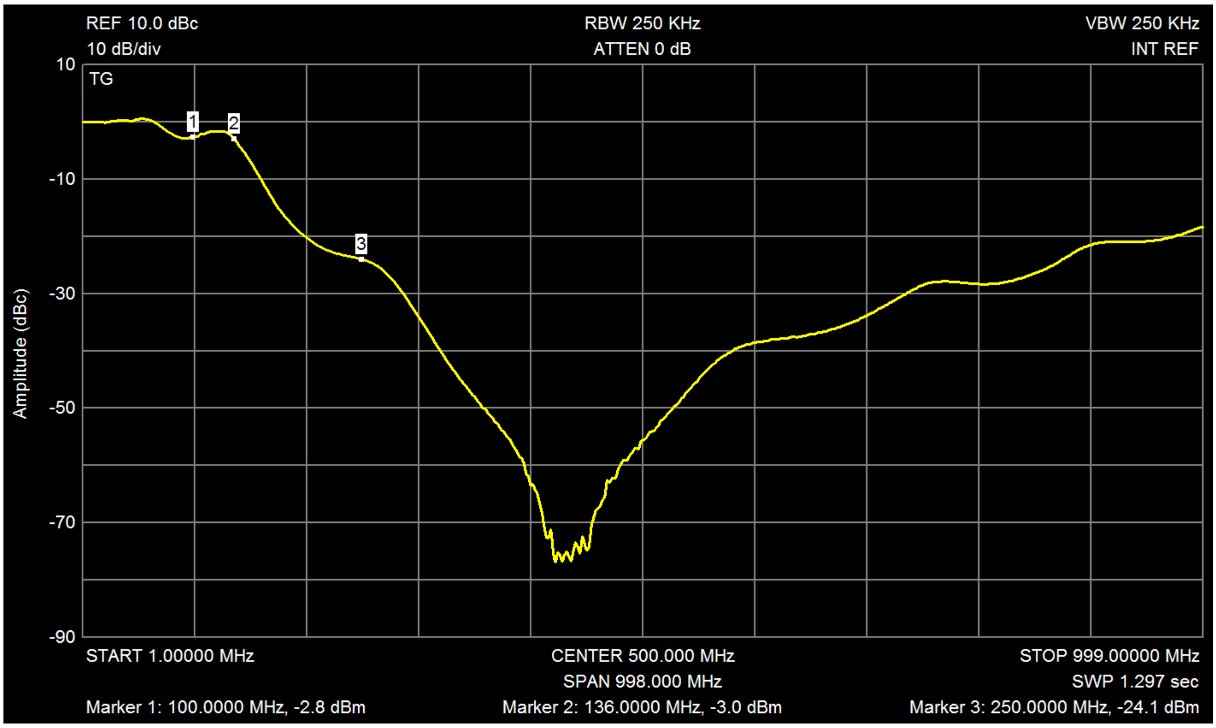

AA-Filter_136MHz

As can be seen from the measurement shown above, the frequency response is far from perfect. We already get 2.8dB attenuation at 100MHz and the actual -3dB corner frequency is 136MHz, but this is still perfectly acceptable for the job. Attenuation at the Nyquist frequency of 250MHz is 24.1dB – far less than what the SDS1104X-E appears to have. If anything, this filter on the SDS1202X-E would yield worse results for the square wave test than the bandwidth limited SDS1104X-E frontend.

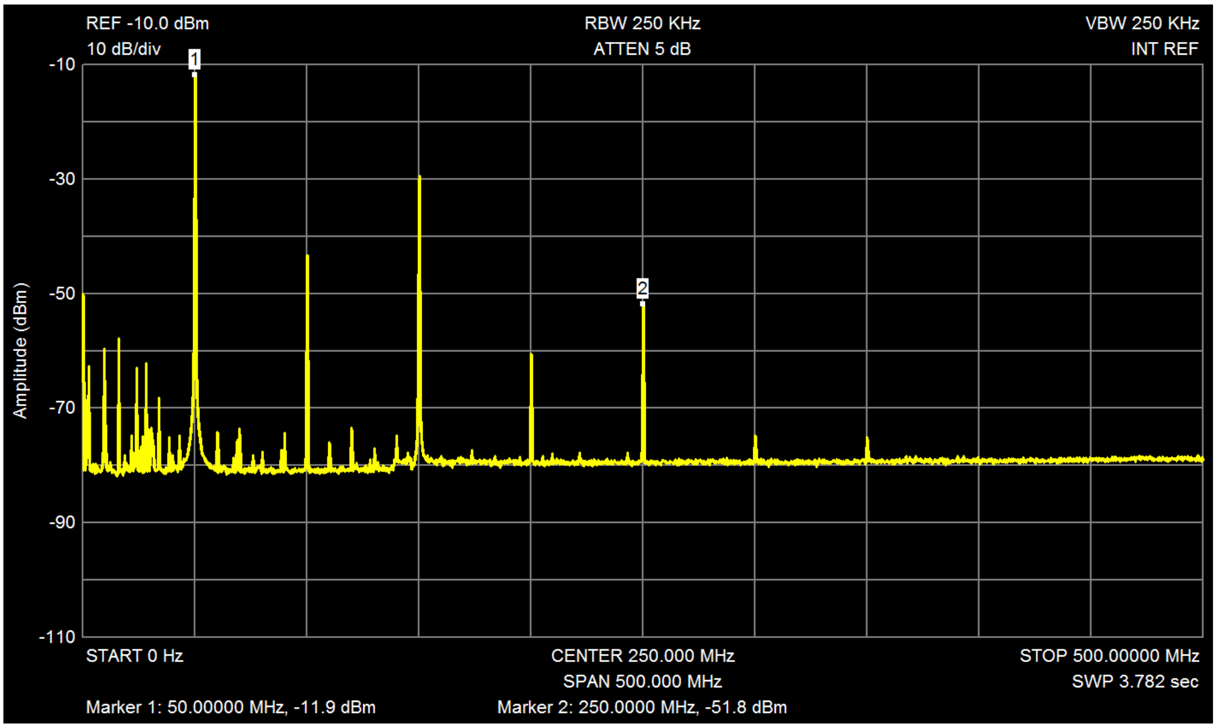

Square_40MHz_Spectrum_filtered_3rd_136MHz

The spectrum of the test signal after the filter is shown above. The strongest harmonic at the Nyquist frequency of 250MHz is now -39.9dBc. As expected, the SDS1202X-E shows an absolutely clean trace with the filtered signal:

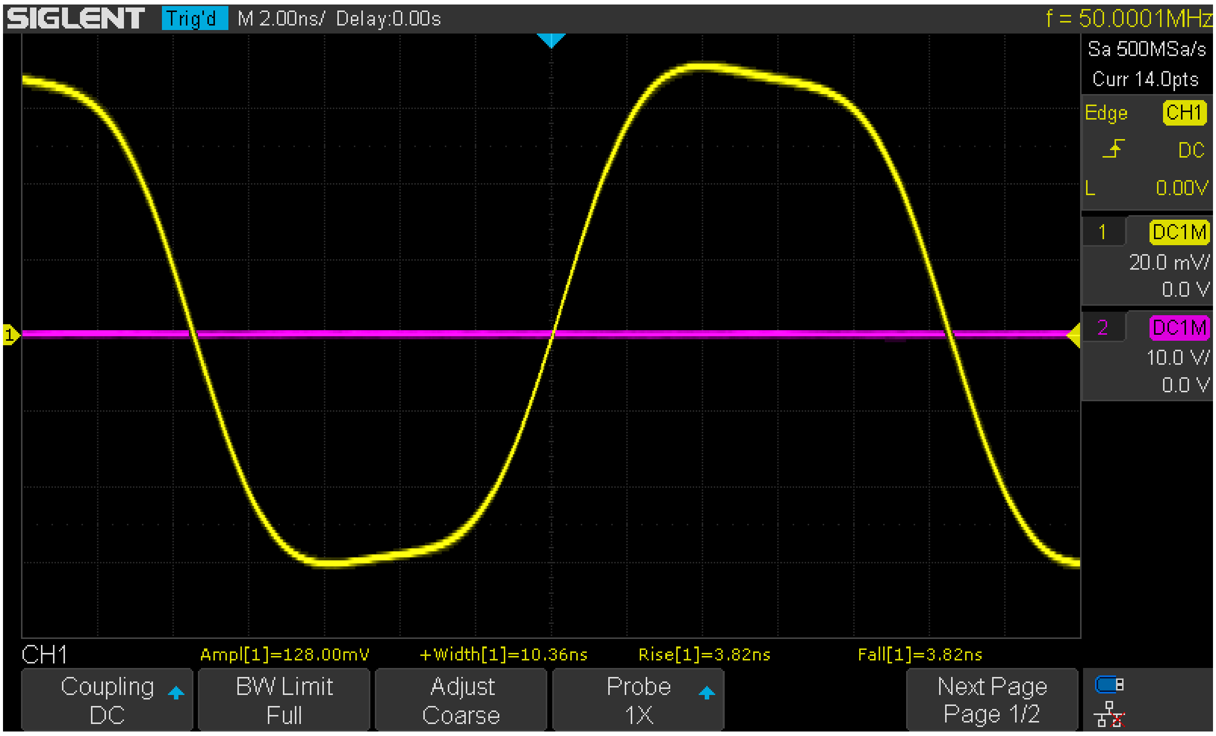

Square_40MHz_BW237MHz_500MSa_Filter_3rd_136MHz

What conclusions can we draw from this?

It looks like the 100MHz SDS1104X-E not only uses essentially the same frontend as the 200MHz models, there is also no physical bandwidth limit fitted. It would have been obvious to just fit an extra capacitor somewhere in order to limit the bandwidth, or have it switchable just like it is the case with the selectable 20MHz bandwidth limit. The venerable Rigol 1052E used such an approach, where the 50MHz bandwidth limit was controlled by the firmware quite simply based on the model number. None of these methods seem to apply here, as the Siglent SDS1104X-E shows aliasing problems that could not occur if it had any true input bandwidth limit. Furthermore, with just a capacitor in the frontend we would expect a 1^{st} order filter slope instead of a 3^{rd} order one. So the actual bandwidth limit has to be somewhere in the digital signal processing and this is of course controlled by the firmware.

The result is somewhat disappointing for those who might have hoped to get a strong protection against aliasing with the bandwidth limited model, but then again, the issue isn't huge and not many users of an entry level scope like this are likely to face signals that fast – signals that normally would require at least a 1GHz scope.

Potential hackers might be disappointed as well, as there is no easy hack by just removing a few capacitors to get the full bandwidth.