Measurement Errors Caused by the Transient Limiter

Tim Williams

Elmac Services
Wareham, United Kingdom
timw@elmac.co.uk

***Abstract*— The standard conducted emissions test may use a transient limiter between the LISN and the measuring instrumentation. This component is non-linear and can introduce significant errors in the measurement, sufficient to completely change the outcome of the test but which are transparent to the test engineer. The different mechanisms by which this may happen are described and modelled, and some mitigation techniques are recommended.**

I. Introduction

The conducted emissions (CE) test applies to virtually every mains powered product, to prove that its RF emissions are safely below the limits established for protecting radio reception from interference coupled via the mains supply.

The test produces reasonably consistent results provided that some precautions are taken. One precaution is intended to protect the safety of the test instrument – spectrum analyzer or measuring receiver. The use of a transient limiter between the LISN and the receiver input prevents high level transients from appearing at the input and potentially damaging it. Its attenuation factor must be corrected for, but otherwise its effect on the test tends to be ignored. Unfortunately there are some circumstances in which the limiter can cause dramatic errors in measurement, certainly enough to cause a failure in a product which should pass, or vice versa. This paper describes these circumstances.

II. The Limiter

In the CE test, the measurement is made of the RF voltage impressed on the mains supply between each phase and earth. This means that the receiver input must be connected to each of the live and neutral lines, via a network which blocks the mains voltage without seriously attenuating the RF signal. This is one function of the LISN (Line Impedance Stabilizing Network). Because it must not attenuate the RF signal, any short transients (a few microseconds or less) on the supply will also be passed through, and if their energy is high enough, they will damage the sensitive front end of the analyzer or receiver. Only a few volts amplitude may be too much, particularly with an unprotected spectrum analyzer.

Any equipment under test (EUT) may produce transients well in excess of this, especially when it is switched off. The di/dt at switch off passes through the LISN's 50µH choke, and the resulting voltage spike is limited only by stray capacitances in the rest of the set-up and may be enough to damage the test instrument. Since a general test lab can never be sure that a given EUT won't create such switching transients by itself, some means is needed of clamping these transients to a safe level without affecting the RF signal being measured. This is the purpose of the limiter [1].

A limiter must therefore present as little attenuation as possible to the RF path throughout the measurement frequency range – for CISPR-based measurements, this is 9kHz to 30MHz – while the signal amplitude remains at the expected level for typical measurements, that is, near or below the limit line for the test. But if the signal amplitude approaches the danger level for the receiver front end, the limiter must present enough attenuation to clamp the amplitude below this level.

The usual way of doing this is to use one or more back-to-back diode pairs across the signal path. In simplest terms, as long as the instantaneous voltage in either polarity doesn't exceed the diode forward threshold voltage (0.6V for silicon) then only a small leakage will flow through the diodes; but as the voltage increases the diode impedance drops and starts to attenuate the signal. The signal source impedance is determined by the EUT and if it were not controlled, would give an unpredictable limiting response. This means that some attenuation must be provided before the diodes. Commercial limiters usually have 10dB attenuation with a flat frequency response, followed by one or more stages of limiting. The simplified schematic of a typical device is shown in Fig.1.

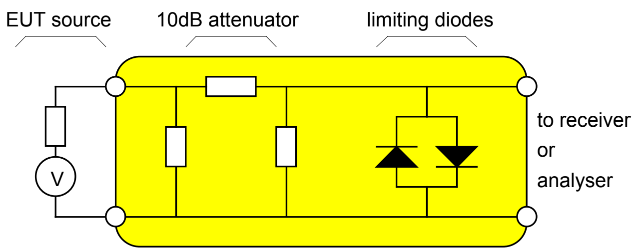

Fig. 1 Simplified limiter schematic

III. Error Mechanisms

This method does indeed protect the input of the receiver or analyzer, and it is widely used in test laboratories for day-to-day CE measurements. Indeed, for many test engineers the limiter is almost an invisible component; the 10dB correction factor is added in the measurement software and it is otherwise assumed to be entirely benign.

For any device in the signal path between LISN and receiver to be benign, it should be entirely passive and linear for all measurement circumstances. This is the case for cables and ordinary attenuators. But with the limiter, if the signal amplitude doesn't remain at the expected level, but exceeds it, then the diode clamp becomes significantly non-linear; and the impact of a non-linear device is either to create harmonics and intermodulation components where none previously existed, or to cause genuine signals to be excessively reduced.

It may be objected that if the signal amplitude at the diodes is high enough to drive them into limiting, then it's surely high enough to breach the emissions limits anyway and so cause a test failure. Leaving aside the issue of accuracy in metrology, this is a reasonable argument for the standard CISPR class A and B limits when the signals that appear over the limits are narrowband and occur within the measurement range; the highest QP limit in these circumstances is 79dBμV, which after the 10dB attenuator is around 600 times lower than the diode clipping voltage. But it fails when:

- the EUT produces broadband interference to be tested against the high-level limits found in CISPR 11 [2];
- the EUT is tested against the usual class A or B limits, but produces high level interference (or even wanted) signals that occur outside the tested frequency range and so are not subject to the limits.

A. Broadband with High Level Limits

CISPR 11 Group 2 Class A limits allow 100dBμV from 150kHz to 500kHz, 86dBμV from 0.5 to 5MHz, and 90 to 70dBμV above 5MHz. If the EUT produces broadband noise that extends right up to this limit across the whole range – not likely, but a theoretical possibility – then it ought to pass the test, but the actual RMS signal amplitude drawn from the LISN would be of the order of 0.9V. Even with a 10dB attenuator before the limiter, this pushes the diodes dangerously close to clipping.

If the noise is pulsed with a low duty cycle, then the quasi-peak detector will drop the indicated value by an amount determined by the pulse repetition rate; so the actual peak amplitude of the signal could be a lot higher than the RMS value quoted above. This is very likely to drive the limiter into clipping; Fig. 2 shows this effect. In so doing, it will flatten the spectrum of the emissions, and therefore possibly reduce the measured signal and hence cause a false pass. But it may also create extra harmonic and intermodulation products, and therefore cause a false failure at the top end of the measurement range.

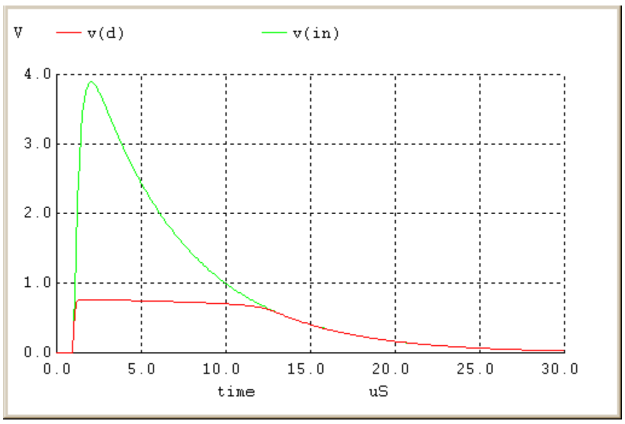

Fig. 2 Clipping a pulsed signal

B. Narrowband, Outside the Tested Range

Most commercial products are measured from 150kHz to 30MHz. Only a few standards call for measurements down to 9kHz and many labs may never test to these standards. An example which demonstrates the problem could be an industrial rack system with a switch-mode power supply of several kW, that must meet the class A CISPR limits. Such a design will probably have an operating frequency somewhere between 30 and 150kHz. As such, say 45kHz, the fundamental and first two harmonics will fall below the tested range. Only the fourth and higher order harmonics (180kHz and above) will be subject to limits. So the design will incorporate a filter which is efficient at and above this frequency, but which need have little or no attenuation below it; and this means that the SMPS emissions at 45, 90 and 135kHz will have a very much higher amplitude than the higher frequencies. Alternatively, the operating frequency may be just below 150kHz, say 140kHz, so that the fundamental is unregulated and the second harmonic is around 280kHz, so that filtering below 280kHz is unnecessary. For a power supply of a few kW, these emissions could well reach several volts, mostly at the fundamental.

Most LISNs are specified to pass signals down to 9kHz even though measurements are only rarely made in this band. So by default, the several volts at 45kHz will pass through to the limiter without significant attenuation. This causes the limiter diodes to clip, generating harmonics of the 45kHz that were not in the original signal at the level passed out of the limiter. These harmonics are measured at the receiver or analyzer *as if they were created in the EUT*, and so result in an unnecessary test failure. The test engineer may well not realise that the low-frequency emissions are high enough to cause such a phenomenon.

The reverse effect can also occur. If the out-of-band signal is high enough, it will drive the diodes into such a low impedance state that they provide a substantial attenuation to the in-band signals that ought to be measured. The result is then a test pass which should be a failure. This is most likely when the out of band signal is at a low frequency, so that its harmonics are not at a high enough level above 150kHz to create a failure themselves.

Naturally, the design of the limiter has a significant bearing on its performance under these overload situations. Schottky diodes, for instance, have about a third of the threshold voltage of ordinary silicon diodes and so will start to clip at about a 10dB lower level. The less attenuation there is before the diodes, also the lower signal level is needed before clipping. So it is possible to find some types of limiter performing better than others, but because of its non-linear nature, any limiter will suffer from the problem at some level.

IV. Reproducing the Effect

It should be understood that none of the above issues are purely hypothetical; they have been observed on real tests, and have been responsible for dramatically different results from different labs, where one lab passed an EUT while another failed it, both with large margins. Until the limiter was isolated as the offending item, there was no obvious explanation for the difference. The limiter itself can be in perfect working order and properly calibrated, and still create the effect. Its calibration will have been done at a low level and a few narrowband spot frequencies, which would not cause any kind of saturation.

C. An artificial EUT

An example EUT has been constructed which demonstrates the effect directly. This injects a 46kHz near-sinusoidal signal onto the neutral line at a level of about 4V p-p (123dBμV rms), enough to cause a degree of clipping in a diode limiter after a 10dB attenuator. Its inherent harmonic distortion is such that all of the harmonics above 150kHz are just, but in some cases only just, below the Class A limit. The lower end emissions measurement is shown in Fig. 3, with a limiter in place, and with it replaced by a passive 10dB attenuator.

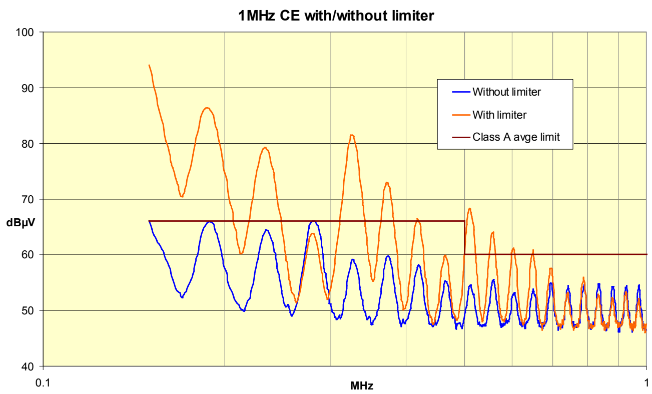

Fig 3 The example EUT from 150kHz to 1MHz, with and without a limiter in place

The limiter creates something like a 20dB increase in the low order harmonics, more than enough to make the product into a clear test failure. The amplitude of the low-frequency signals can be seen in Fig. 4 which shows the spectrum between 25kHz and 200kHz. For a standard CISPR 22 test, the signals below 150kHz would never be investigated.

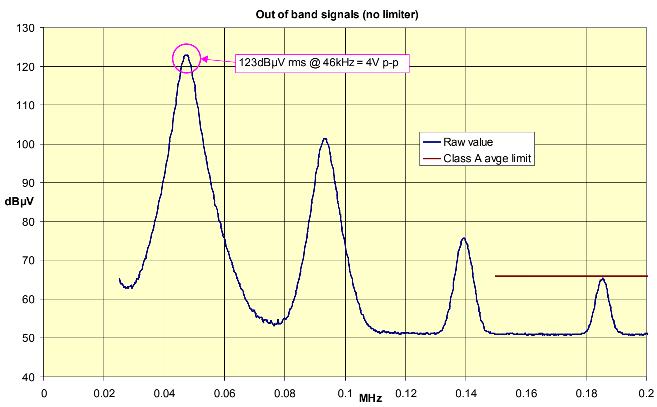

Fig. 4 The lower frequency fundamental and harmonics

D. Modelling the Mechanism

The schematic of a typical limiter has been modelled in WinSpice [3] and this provides an opportunity to derive the level of low frequency sinusoidal signal which will just cause the limiter to create a false test failure.

Two circuits have been considered; one is a simple single diode pair constructed experimentally, the other is a two diode pair commercial unit [4]. In both cases their circuits were modelled with a pure 140kHz sinusoid input (just below the CISPR band B edge) from a source impedance of 50 ohms and with a load impedance also of 50 ohms. This gives a theoretical attenuation of around 15.5dB with the component values used. The output signal was analyzed with the WinSpice Fourier transform function to give the output amplitudes of the fundamental, third, fifth and seventh harmonics, versus input level (the even harmonics are generated only by diode asymmetry and need not be analyzed). The results are shown in Fig. 5 and Fig. 6.

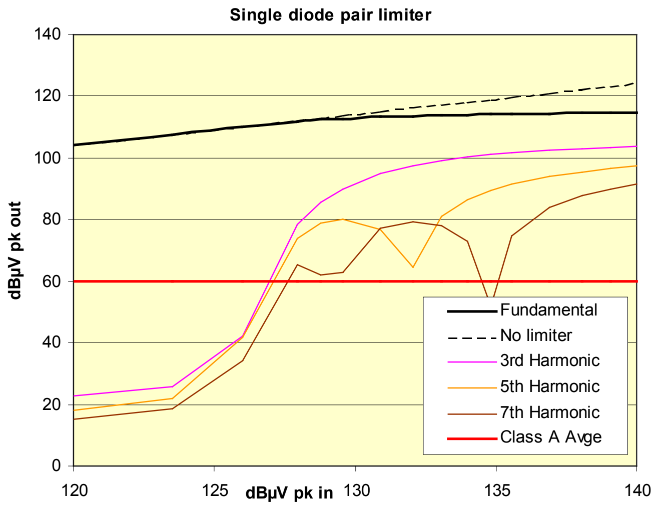

Figure 5 Harmonic output vs. input of the single pair limiter

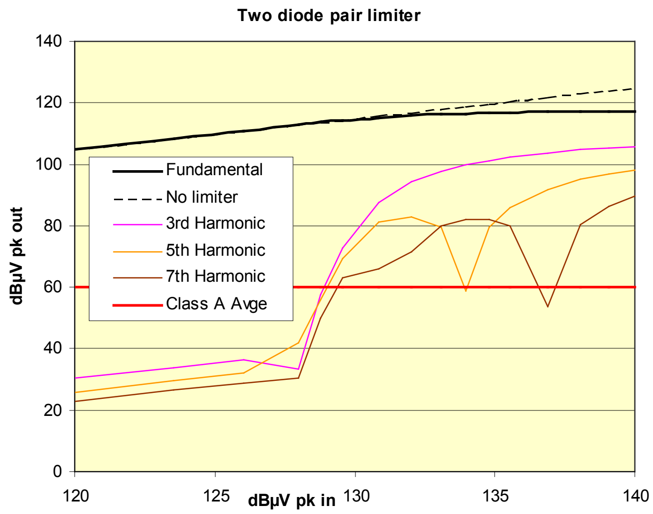

Figure 6 Harmonic output vs. input of the dual pair limiter

The two regions of operation can be clearly seen in these graphs; at low levels the harmonic distortion is well below the level that would exceed the limit, but above around 2V peak the harmonic content climbs rapidly until the diodes are saturating. The dotted black line in each case shows the level of 140kHz fundamental that would be seen at the output if no limiter, only the equivalent attenuation, were present. The compression introduced by the limiter above its threshold is evident.

The exact input level at which the harmonic content exceeds a given limit depends on the limiter construction. For the two-pair commercial device it is 2.7V peak, for the single-pair unit it is 2.25V peak. This will give a third harmonic around 60dBμV; the higher harmonics reach the same level at almost the same input voltage. A similar analysis could be done for other limiter circuits, limit levels and input frequencies, but the actual threshold value will not be dramatically different as long as the same types of diodes are used. Of course, this analysis has been done for a pure single-frequency source at 140kHz, whereas real sources will have harmonic content of their own so that the extra distortion introduced by the limiter will become significant at a slightly lower voltage.

V. Mitigation

Once the phenomenon and its effect is understood, there are a few actions that can be taken to improve the reliability of the measurement.

E. No Limiter

The first and obvious solution is not to use a limiter at all. This requires the measurement receiver to be robust enough to cope with all likely over-voltages, and for many test labs it is the preferred solution since this source of error is permanently removed from the test. It's not advised for those labs which use spectrum analyzers for pre-compliance purposes, since a spectrum analyzer is inherently unprotected and can easily be damaged by unexpected switching spikes.

F. Temporary removal of limiter

Another approach is to remove the limiter temporarily, and replace it with a passive 10dB pad. While this is simple and effective – if there is a change in the measured spectrum, the limiter is creating its own extra signals, or blocking the true ones – it may also be inconvenient if the limiter is permanently wired in to the system. Since the whole purpose of the limiter is to protect the receiver, when it is removed the protection is lost. This is usually acceptable if it is done temporarily, ensuring that the EUT continues operating without creating switching spikes, but it is not a permanent solution. It could easily be achieved if the limiter included a non-latching switch to remove the diodes from the circuit.

G. High pass Filter

A different solution would be to place a passive high pass filter with a cut-off just below 150kHz, after the LISN but before the limiter. This should attenuate the un-measured signals to the level at which they do not cause clipping, without affecting the measured frequency range. This solution does not address the problem of clipping on broadband pulsed noise.

H. Extra attenuation

A final possibility would be to place extra passive attenuation between the limiter and the LISN so that the signal at the limiter is reduced. This must be corrected for in the transducer factors, and it reduces the sensitivity of the system so cannot be taken very far, but it allows for a "reality check" like the first option without defeating the protection offered by the limiter. Because diode clipping is inherently a non-linear effect, a reduction of 10dB in the signals at the input will result in a greater than 10dB reduction in the output signals that are due solely to the clipping. Of course, switching in the receiver or analyzer's input attenuator will not be useful, because this comes into play after the limiter.

VI. Conclusions

If the conducted emissions test produces unexpected results, suspect the limiter. This paper has described a number of mechanisms by which the non-linearity of the limiter can falsify the measurement. Even if the results are apparently acceptable, still suspect the limiter if there is any concern that the signals produced by the EUT may be at a sufficiently high level to drive it into non-linearity. There are a number of methods by which any errors due to the limiter can be isolated. If the test receiver in use is sufficiently well protected not to need a limiter, don't use one "just in case" – it is a further source of potential error which a test lab can do without.

Acknowledgment

The author is grateful to Bernard Wragge-Morley and Peter Metcalfe for their comments, and to John Melling for initially drawing attention to the problem.

References

[1] T. Williams, *EMC for Product Designers*, 4th Edition, Oxford: Elsevier, 2007, section 6.2.2
[2] *Industrial, scientific and medical (ISM) radio-frequency equipment — Radio disturbance characteristics — Limits and methods of measurement,* CISPR11: 1997
[3] M. Smith, WinSpice3, http://www.winspice.com
[4] *Operation and Service Manual*, 11947A Transient Limiter, Agilent Technologies, 1991