The hazard of the transient limiter

Tim Williams, Elmac Services

Introduction

The conducted emissions (CE) test is ubiquitous: virtually every mains powered product has to undergo such a measurement, to prove that its RF emissions are safely below the limits established for protecting radio reception from interference coupled via the mains supply. Even DC-powered apparatus has to take the CE test for some applications.

Generally, it's a simple test to do and produces reasonably consistent results provided that some precautions are taken. One of these precautions is not to do with protecting the quality of the test, but of protecting the safety of the test instrument – spectrum analyzer or measuring receiver. It's the use of a limiter between the LISN and the receiver input. This prevents high level transients appearing at the input and damaging it, and once it is in place and its attenuation factor is corrected for, it tends to be forgotten. But 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. How does this occur, and what can be done about it?

The limiter

In the CE test, the measurement is made of the RF voltage impressed on the mains supply: in the case of a single phase supply, between the live and earth, and between neutral and earth. This means that the receiver input must be connected to each of the live and neutral lines, not directly of course, but 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 the RF signal must be passed through, 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 enough for this, particularly with an unprotected spectrum analyzer (receiver inputs are usually better protected and more robust).

Any equipment under test (EUT) may produce transients well in excess of this, especially when it is switched off and hence the current through the LISN's choke is interrupted. The di/dt at switch off is theoretically infinite, and since the usual commercial LISN includes a 50µH choke, the voltage transient given by $V = -L \cdot di/dt$ is also theoretically infinite – in practice, limited by stray capacitances in the rest of the set-up. Certainly it may be enough to damage the test instrument. Since, in a general test lab, you 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.

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. 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. This should only happen on transients, not continuously, so that the normal measurement accuracy is unaffected.

The most typical way of doing this is to use a plain back-to-back diode pair 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. In order to work properly, the signal source impedance must be controlled, and this means that some attenuation must be provided before the limiter. Commercial limiters usually have 10dB attenuation with a flat frequency response, followed by one or more stages of limiting. During the test this is included as part of the transducer factors, and means that the measurement is 10dB less sensitive than it could have been, which is not normally a problem with CE measurements. The simplified schematic of a typical device is shown in Figure 1.

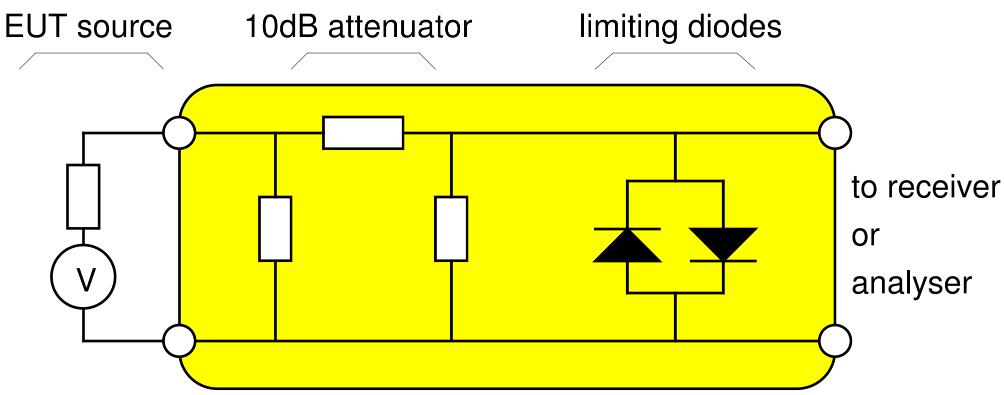

Figure 1 Typical limiter schematic

The hazard

This method works well and 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 that's the end of it. Unfortunately, it isn't.

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. Notice the condition that was quoted above: "present as little attenuation as possible while the signal amplitude remains at the expected level for typical measurements". 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 signals that need to be measured to be excessively attenuated.

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, so what's the problem? 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. It becomes less tenable under the following circumstances:

- when the EUT produces broadband interference that will be tested against the high-level limits found in CISPR 11;
- when 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.

Let's discuss each of these cases.

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. Given a 10dB attenuator before the limiter, this does not actually push the diodes clearly into clipping, but it certainly brings the level close enough that the limiter should not be regarded simply as a linear, benign, device.

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 frequency; so the actual peak amplitude of the signal will be potentially a lot higher than the RMS value quoted above. This is very likely to drive the limiter into clipping. In so doing, it will flatten the spectrum of the emissions at their highest level, and therefore possibly reduce the measured signal and hence cause a false pass; but also it may create harmonic and intermodulation products at higher frequencies than were present in the original, and therefore cause a false failure at the top end of the measurement range. Figure 2 shows this effect.

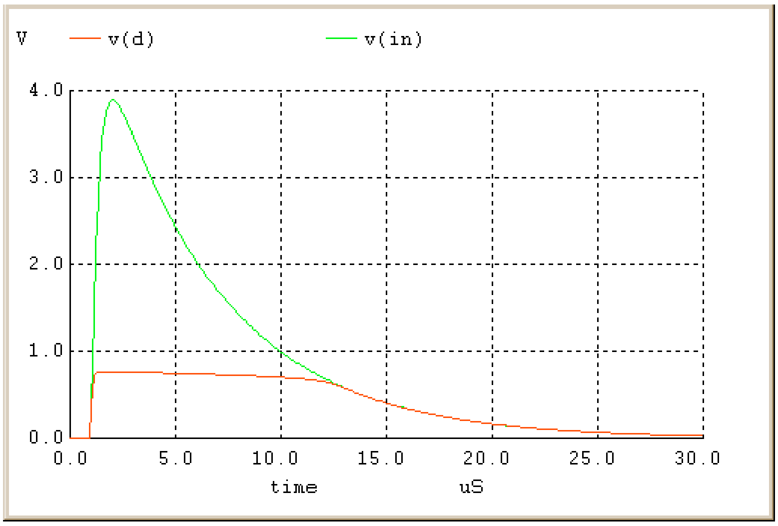

Figure 2 Clipping a pulsed signal

Narrowband, outside the tested range

The conventional CISPR CE frequency range that most products have to be tested to is from 150kHz to 30MHz. Measurements down to 9kHz are called for in only a few standards and many labs may never test to these standards. Imagine, then, that you have to design a large telecom rack product which includes a switch-mode power supply, that must meet the class A CISPR limits. The power supply may be rated at several kW. Such a design will probably have an operating frequency somewhere between 20 and 50kHz. 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 you will design 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 be very much higher amplitude than the higher frequencies. For a power supply of a few kW, these emissions could well reach several volts.

Because of the CISPR Band A (9-150kHz) specification, most CE LISNs will be 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 never in the original signal – or at least, never in it 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, where they result in an unnecessary test failure. If you are unaware of this mechanism, and never make measurements below 150kHz, you 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.

An example

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. Even the same lab could get different results on the same EUT, if it had different CE measurement set-ups. Until the limiter was isolated as the offending item, there was no obvious explanation for the difference. Remember that 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.

An example EUT has been constructed which demonstrates the effect directly. This simply 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. The harmonic distortion of the 46kHz signal is such that all of the harmonics above 150kHz are just, but in some cases only just, below the Class A limit. Whilst this has been artificially created as an example, it is by no means an unrealistic scenario for many types of product as discussed above, where the filtering has been tailored to the minimum necessary to stay legal. The lower end emissions measurement is shown in Figure 3, with a limiter in place, and with it replaced by a passive 10dB attenuator.

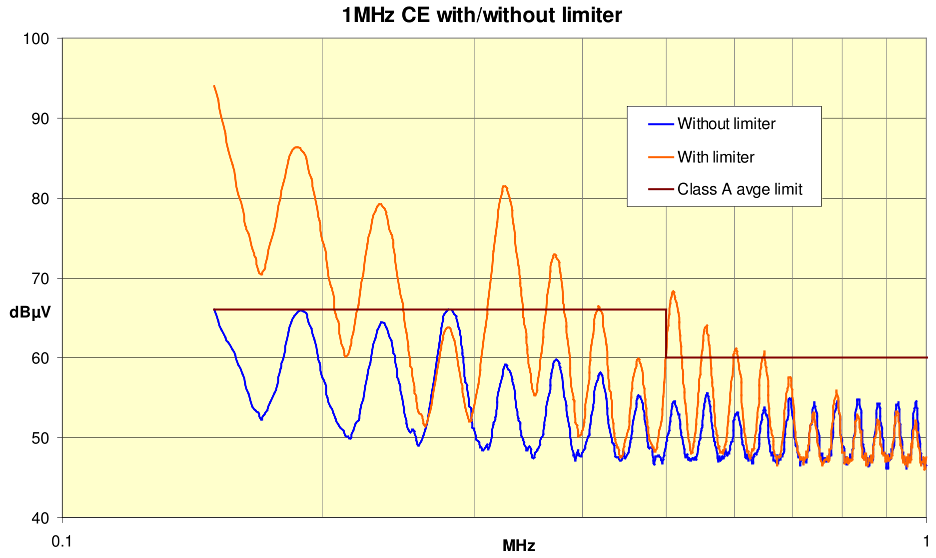

Figure 3 The example EUT from 150kHz to 1MHz, with and without a limiter in place

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

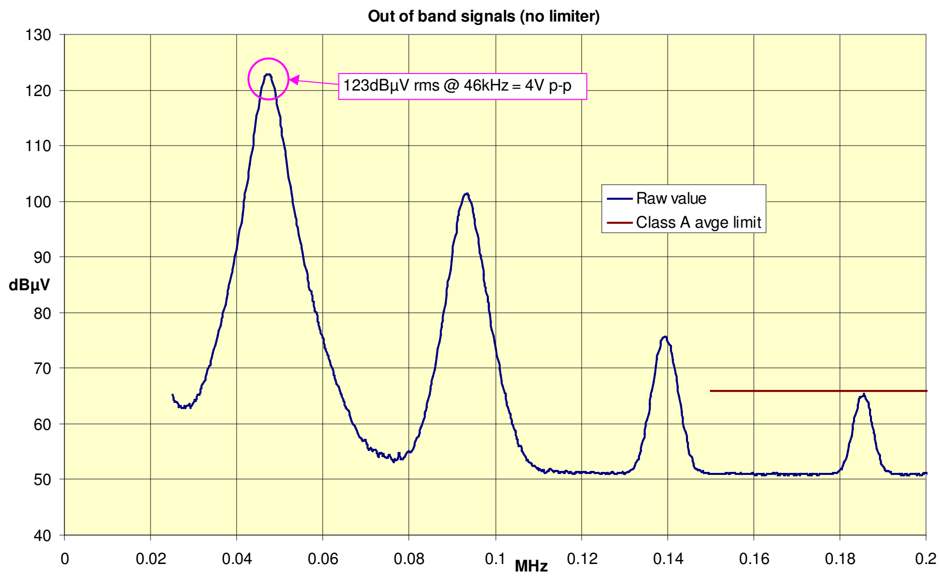

Figure 4 The lower frequency fundamental and harmonics

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.

The first and obvious solution is not to use a limiter at all. This is fine if your measurement receiver is robust enough to cope with all likely over-voltages, and for many test labs it is the preferred solution since one source of unexpected 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 switching transients created within the EUT.

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 cumbersome if the limiter is permanently wired in to the system. Since the whole purpose of the limiter is to protect the receiver, when you take it out you also lose the protection. This is usually acceptable if it is done temporarily, ensuring that the EUT continues operating without creating switching spikes, but it's not a permanent solution. It would be easy to achieve if limiter manufacturers were to put in a "bypass" switch that removed the diodes from the circuit, but they don't.

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. But such filters are not widely available, although easy enough to construct, and this solution does not address the problem of clipping on broadband pulsed noise.

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 of course be corrected for in the transducer factors, and it reduces the sensitivity of the system so cannot be taken very far, but it would allow 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.

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.

Elmac Services
Wareham, Dorset
UK
Tel +44 1929 558279
e-mail timw@elmac.co.uk web www.elmac.co.uk

This is version 1.1 of www.elmac.co.uk/pdfs/limiter_hazard.pdf