Another EMC resource
from EMC Standards

Suppressing emissions by using data scrambling and spread spectrum hardware design techniques

Helping you solve your EMC problems

9 Bracken View, Brocton, Stafford ST17 0TF T:+44 (0) 1785 660247 E:info@emcstandards.co.uk

Suppressing emissions by using data scrambling and spread-spectrum hardware design techniques

Keith Armstrong, www.cherryclough.com

First published in the EMC Journal, Issue 109, November 2013
(Modified 27 Nov 2013 after reading [21])

Contents

Before I get to the 'meat' of this article, a little background is required.

1 EMC should not be left to EMC engineers

EMC engineering is often assumed to be concerned with mitigating EMI by using techniques such as filtering, shielding, galvanic isolation, 'grounding' (so-called, although it has nothing to do with the soil, or with the green/yellow striped wire in the mains power cord), surge and transient suppressors, etc. But cost-effective EMC engineering is so much more than this!

These mitigation techniques are typically used by EMC test engineers when a customer's product fails an emission or immunity test, and they are asked to help fix the non-compliance. The overall result is generally an increase in BOM cost; weight and size, caused by adding parts that have nothing (directly) to do with functionality and are (incorrectly) seen as being the necessary price of complying with EMC regulations.

Project planning that uses the traditional approach of '*achieve the specified functionality and then throw EMI mitigation at it until it passes its EMC compliance tests*' is guaranteed to delay time-to-market by an unpredictable amount, add cost, weight, etc., and also runs the risk of having to make major changes to the mechanical design to fit the required filters, shielding, etc.

Manufacturers who use this approach to projects inevitably suffer unpredictable financial risks, up to and including complete failure of their company. This is not because of the effect of the increased BOM cost on profitability – as is usually believed (because it once used to be true, some decades ago, when products had sales lives of 10 or even 20 years) – but because of the delays in time-to-market, see [1] and [2].

In common with all of the experienced EMC consultants I have ever met, I strongly recommend using good EMC engineering design practices from the 'blank sheet of paper' start of any project. I suppose that these days we ought to call this the 'blank CAD screen' project stage, instead.

These good EMC engineering design techniques need to be implemented by mechanical, hardware, FPGA, software and/or PCB engineers, who project managers often seem to assume have nothing to do with EMC, with the results summarised in Figure 1.

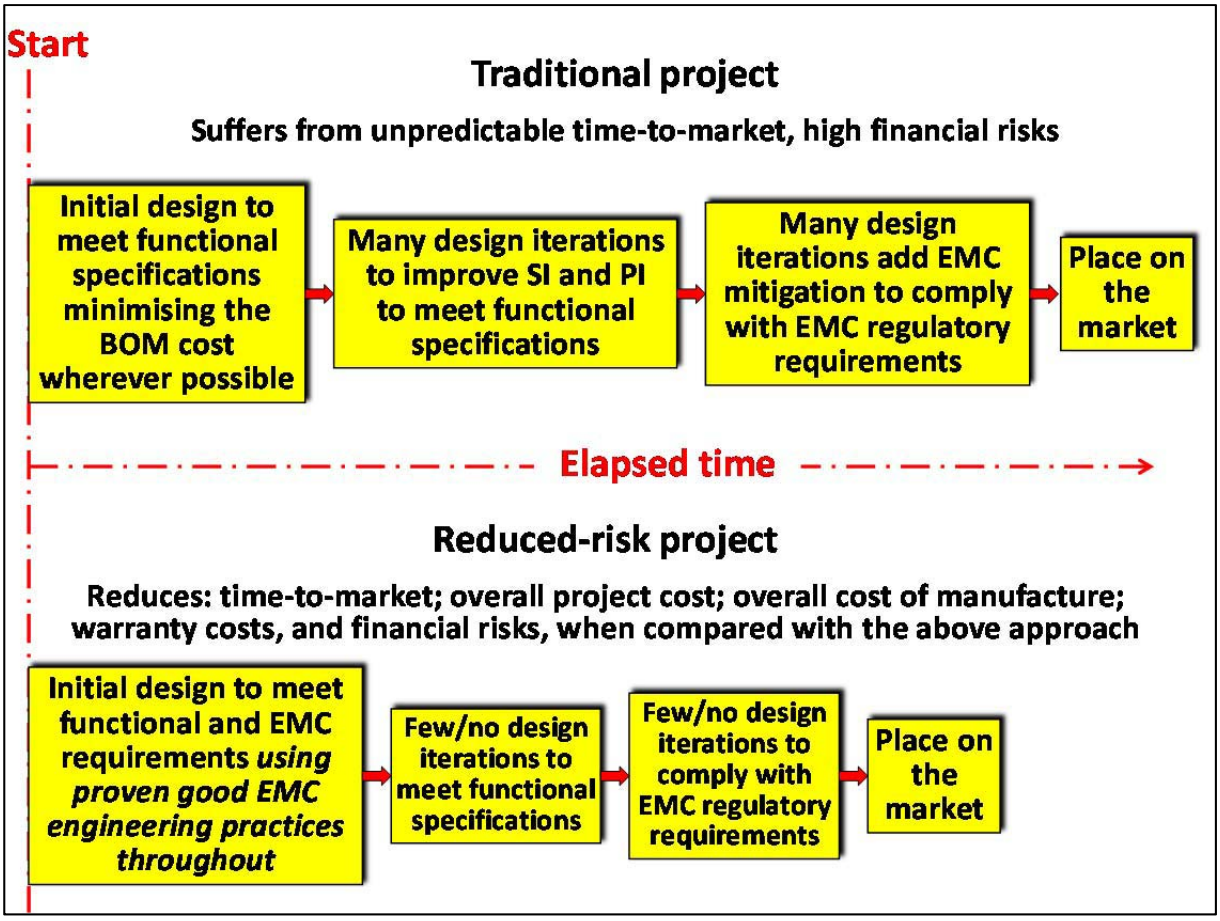

Figure 1 Comparison of traditional and reduced-risk projects

Using good EMC engineering practices from the start of a project automatically takes care of most signal integrity (SI) and power integrity (PI) issues whilst it sets about achieving the most cost-effective design that complies with EMC regulations. The cost, weight and size of the EMI mitigation required (filters, shielding, etc.) will be the least – and there will be no unpleasant surprises in the EMC test lab.

Design and development project costs and timescales will be reduced because fewer design iterations will be needed for both hardware and software to achieve the functional specifications and reliability. (In fact, experience over the last 23 years has shown that functional specifications can even be improved upon, without additional cost or delay.)

Warranty returns are reduced because of the inherent EMI robustness of the design, something that is not ensured by passing the standard EMC tests for immunity. The overall result is that a less costly design/development project gets the new product to market sooner, a very important goal these days, see [2] – whilst reducing its overall cost of manufacture (which is more important than the BOM cost, see [1]).

Companies who have fully adopted my recommendations for using good EMC engineering design from the start of a project have often told me of the great results they have received, which their CEOs appreciate as reduced financial risks with increased profitability. One company that has recently fully adopted my EMC design recommendations is documenting the results in detail, and I hope to publish a joint article with them on this in the near future.

Regular readers of my articles in the EMC Journal [3] and other EMC publications; people who have read my books [4] and those who have attended my training courses [5] should already be familiar with my approach, which I have been applying with great success for 23 years now.

The subject of this article is one of the good EMC engineering design techniques that I recommend should be applied from the start of any electronic project. Like most such techniques it adds a little design complexity, and sometimes adds a little to the BOM cost, but helps to achieve huge savings in time-to-market and profitability.

Like so many such EMC design techniques it is easy for hardware/software/PCB designers to apply early in a project, but very difficult and costly indeed to apply when the design is supposed to be finished but the product is failing its EMC compliance tests.

OK, that's enough background.

2 Data Scrambling techniques

When we are testing products for emissions, we usually find it is the narrow peaks in the emitted noise spectrum – often called 'spikes' – that exceed the limit lines and cause us compliance problems.

Continuous RF immunity tests generally fail at exactly the same spike frequencies, so techniques that reduce or eliminate spikes from the emissions spectrum also help pass RF immunity tests.

In the EMC world, we are quite familiar with one way of dealing with this – spread-spectrum clocking – which replaces fixed frequency clocks with ones that are frequency-modulated with a much lower (typically audio) frequency.

But spread-spectrum clocking is just one of a number of techniques which can be used to spread a signal's power over a wide range of the frequency spectrum, and reduce the power of any particular frequency component (i.e. a peak or spike) in the original signal [6].

I shall discuss 'Data Scrambling' in this section, and 'Spread-Spectrum' clocking' afterwards.

Data scrambling can be done in hardware and/or software, but is most often applied to high-speed serial data and so is done in hardware. Data that has been perfectly 'scrambled' would have a spectrum that was close to being 'white noise', with no visible peaks or spikes at any frequencies.

Datacomms designers sometimes refer to scrambling as one of several techniques under the general heading of 'Spread Spectrum' techniques. This is strictly correct, but I've never heard this usage in the EMC community – where the term 'spread-spectrum' seems to be universally used to refer only to the specific technique of frequency-modulating a clock.

When cost-effective design and manufacture is an important issue, as it is for most manufacturers in competitive industries, it is often found that cables and their connectors are the cause of many EM emissions problems.

Because of this, scrambling of data in cables has a very attractive cost/benefit ratio. It is used to great effect in, for example, 100/1000BASE-T Ethernet, which I will now briefly describe as an example of a very effective design approach.

100/1000BASE-T Ethernet communications use asynchronous signalling to avoid the need for a separate conductor for the clock. This may have been done to reduce costs, but is good for EMC because an additional clock line would need good-quality shielding and well-shielded connectors to prevent it's

harmonics from causing radiated emissions above the CISPR 22 or CISPR 11 (= EN 55022 or EN 55011) limits – and/or would need to use spread-spectrum clocking (see below).

100/1000BASE-T also uses a handful of logic gates to scramble the asynchronous serial data, applying a deterministic algorithm (which a handful of gates in the receiver can then unscramble) to prevent any dominant frequencies from appearing in the signal spectrum – for example caused by repetitive data sequences such as 101010101010.

This was done specifically to help meet EMC emissions regulations whilst using low-cost 'physical layer' hardware such as unshielded cabling.

Scrambling the data means that no single frequency is sent for any significant period of time, spreading the signal power over a range of frequencies. This effectively 'whitens' the frequency content of the signal, making it appear more like random noise, which has a broadband spectrum and prevents any particular frequency component from having a higher power than any other [6]. It also has the advantage of using more of the available bandwidth, allowing higher data rates to be achieved without increasing the upper frequency of circuit operation or emissions.

[7] says that 1000BASE-T:

- Uses algorithmic mapping and inverse mapping from octet data to a quartet of quinary symbols and back (the PAM-5 modulation, with the 5^{th} bit used for error correction).
- Has uncorrelated symbols in the transmitted symbol stream.
- Has no correlation between symbol streams traveling both directions, on any cable pair combination (1000BASE-T uses cabling that has four twisted pairs, called A, B, C and D).
- Has no correlation between the symbol streams on pairs A, B, C and D.

1000BASE-T also uses 'Trellis Coded Modulation', to make its 5-level signalling ('PAM-5' modulation) as robust as 100BASE-TX, which uses 3-level signalling. This is implemented as a two-step approach:

- Convolutional Encoding to convert scrambled data octets to 10-bit words (so-called 8b/10b encoding).
- Mapping by Set Partitioning to get 6dB noise immunity gain.

We might expect the 8b/10b encoding to increase the data rates and thereby increase emissions, but [8] tells us that the reverse is the case because the 8b/10b algorithm used creates a special 10-bit sequence that minimises the number of 0-to-1 transitions, and so actually reduces the emissions. 8-bits of non-encoded data requires 7 transitions, whereas using the 10-bit Trellis modulation algorithm reduces this to 3.

Figure 2 shows the resulting spectrum of 1000BASE-T Ethernet signals, according to [9]. Notice that there are no narrow 'spikes' in the spectrum caused by clock harmonics, and no 'bumps' caused by data (e.g. caused by the NRZ (see [10]) data pattern causing strong sidebands of the data clock frequency and its harmonics). Figure 2 also shows the 'eye pattern' of the PAM-5 modulation that transmits 5 bits at the same time.

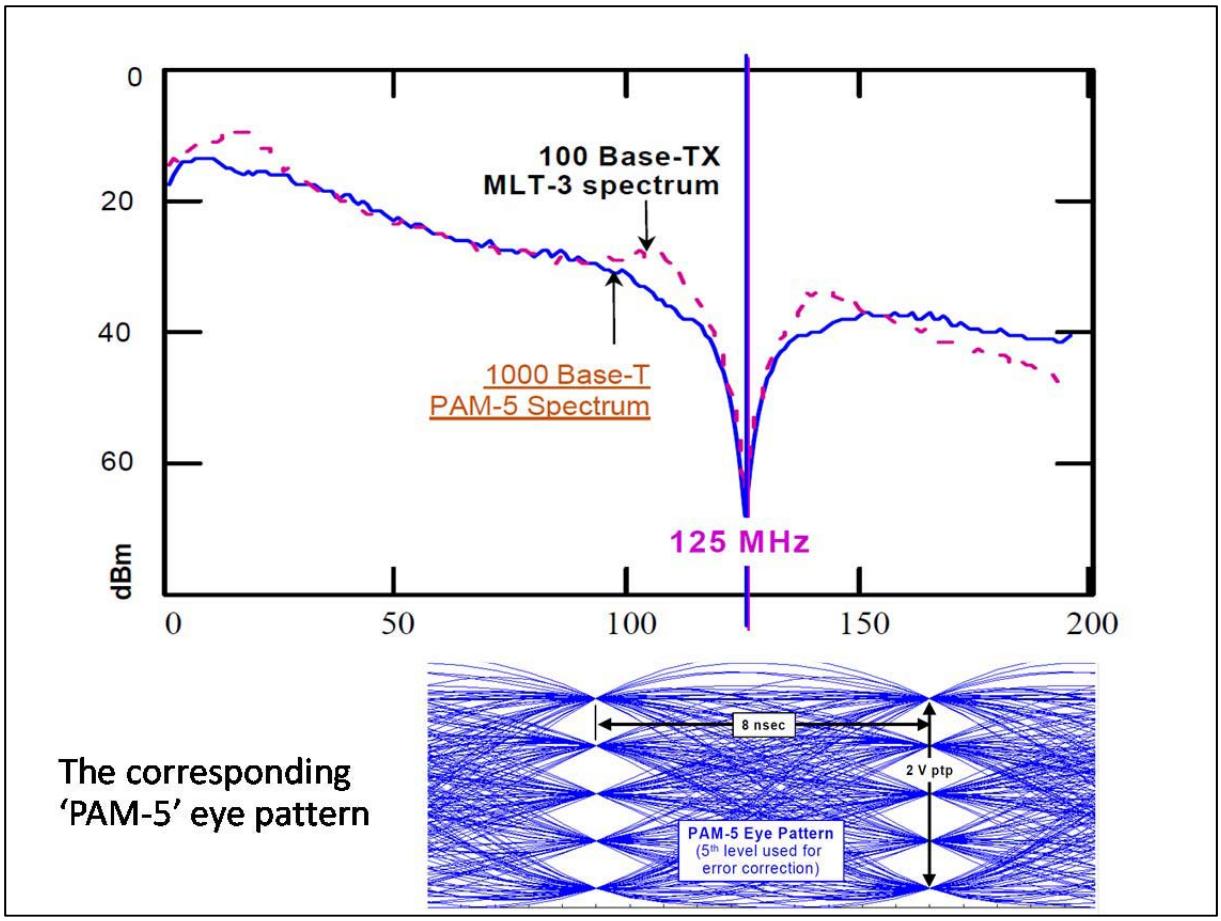

Figure 2 The spectrum of the 1000BASE-T Ethernet signals (from [9])

In this article I am using some figures that show the spectrum of certain signals, and other figures that show the emissions spectrum from certain EMC tests.

These two types of spectrums are not directly comparable, except that we know that a 'non-spiky' signal spectrum will not give rise to a spiky emissions spectrum – so these two types of spectrum graph are comparable in this regard.

Note that even perfectly scrambled data may still suffer from spikes caused by common-mode noise on the PCB, plus peaks and dips in its emissions spectrum due to a variety of systematic resonances in PCBs, enclosures and cables. These should be prevented by using other good EMC design techniques from the start of a project, see [3] [4] and [5].

When Quasi-Peak detectors are used in emissions tests, clock harmonic 'spikes' usually measure up to 1dB less when measured with a Quasi-peak detector, than when they are measured with a Peak detector. At the other extreme, true random noise measures 15dB less Quasi-peak than Peak.

So a measure of the effectiveness of a data scrambling design is by how much a Peak detector spectrum reduces when a Quasi-peak measurement is made on it. It might prove difficult to achieve the full 15dB reduction when measuring Q-peak without heroic scrambling design, but 5 to 10dB reduction should be readily achievable.

Scrambling all of the data that is communicated between ICs would be ideal, but to use scrambled data for parallel data communications between integrated circuits (ICs) on a printed circuit board (PCB) would probably cost more overall than using a BLS (board level shield) to reduce emissions.

But where ICs on the same PCB communicate using high-speed serial data streams (usually using differential transmission line trace pairs), scrambling the data using techniques similar to above, using a

couple of handfuls of spare gates in a modern FPGA, could well be a cost-effective way to reduce the emissions caused by the many sources of differential skew that such transmission-lines suffer from, perhaps avoiding the need to replace the FR4 dielectric with a more expensive one.

I have seen products where an especially powerful *component* operating with high rates of dV/dt was responsible for the worst spikes in the emissions spectrum, and would have significantly benefitted from having scrambling or spread-spectrum techniques applied to the signals or power-switching that it was handling.

In one particular instance of an excessively noisy gigabit laser diode, when I suggested data scrambling as a solution I was told that it had been discussed by the design team at the start of the project, but they hadn't realised how important it would be and so decided to use unscrambled data. Now, at what had been supposed to be the end of the project, they were faced with a small number of very costly and time-consuming options if they wanted to sell a legal product.

Using scrambled data would not have added much time or cost to the original design, and in my view would have been well-worth doing 'just in case'. I often see examples of powerful and low-cost good EMC technique being ignored because of misplaced optimism about the likely EMC test performance of the final product design.

Some data scrambling techniques also provide benefits for EMC immunity, and yet a further reason is to make more efficient use of the available bandwidth, up to the point where an entire band may be utilized for communication, as we are increasingly seeing with digital communications in general.

It's not just product designers who should consider the use of data scrambling techniques. Some industry groups working on common technical standards for standardised interfaces often choose poor data scrambling techniques (or none) – locking everyone who wants to use that interface into an EMC design that is *always* going to be difficult and/or costly to make compliant.

However, industry groups making bad EMC design decisions is nothing new. I have seen an industry standard for parallel datacomms that *required* the use of flat ribbon cables carrying 50 digital signals with just one common return conductor (and that along one edge, not even in the middle!).

I have also seen a pro-audio industry standard for 125Mb/s serial datacomm's that specified the use of 75 Ohm video coax originally used for analogue signals at up to 6MHz simply because there was a large installed base of such cables in that industry's buildings. Both of these examples created huge EMC problems for anyone whose products needed to use these standard interfaces.

For example, Figure 3 (taken from [11]) shows an example of a radiated emissions test on an HDMI system, revealing narrow 'spikes' in the spectrum caused by clock harmonics. The upper graph was obtained using a good quality well-shielded HDMI cable, whilst the lower graph shows the effect of adding a ferrite common-mode clamp to the cable.

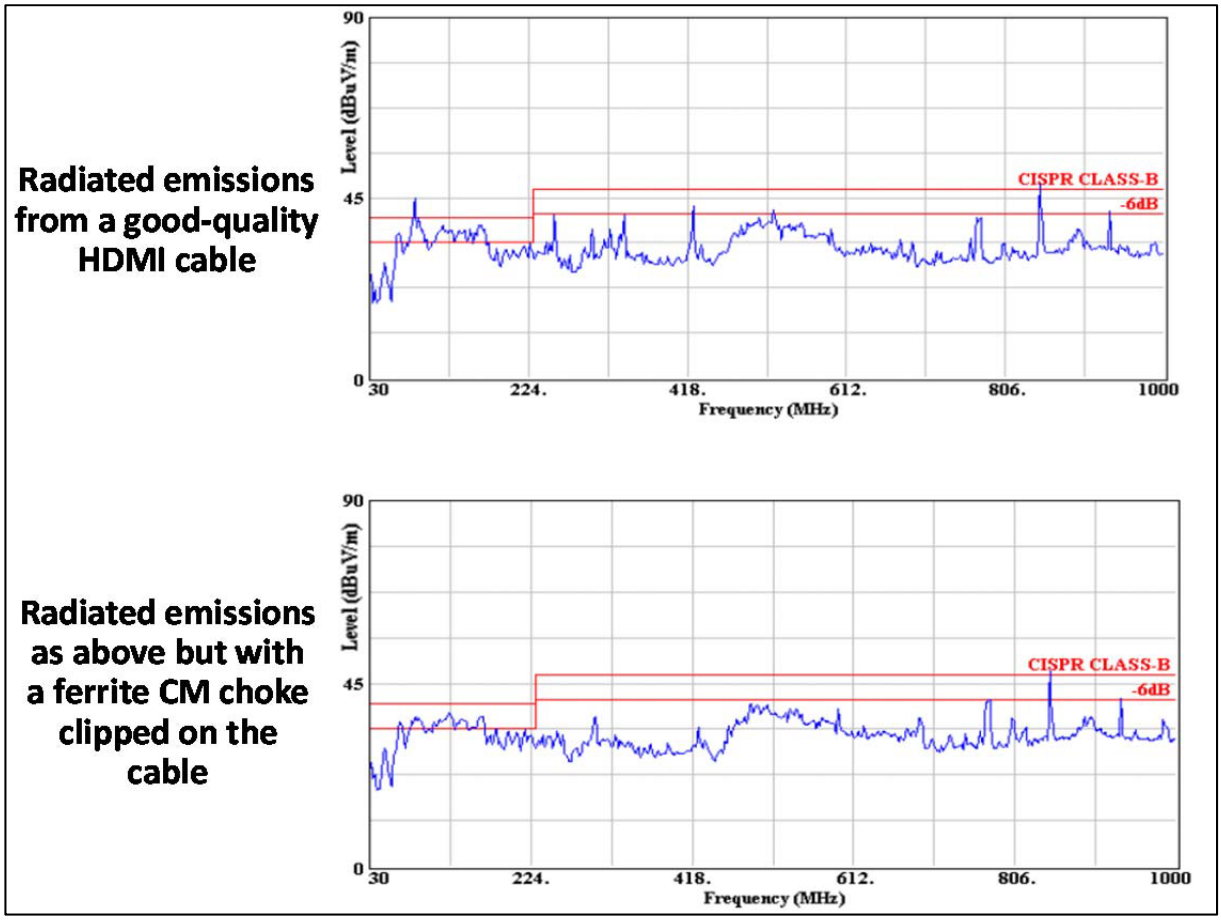

Figure 3 Example of HDMI radiated emissions

As I said above, it is usually such narrow spikes that cause us problems for regulatory EMC compliance, partly because they don't reduce much (if at all) when measured Quasi-peak.

Without a well-shielded HDMI cable, there is no chance of meeting even the 'Class A' CISPR 22 emissions limits. Adding a few gates in transmitter and receiver to scramble the data allows Ethernet to use less-costly cabling, complying with Class A even when using UTP (unshielded twisted-pair) cables and less costly connectors. This is why although they might carry the same data rates, an HDMI interface will need substantially more costly connectors and cables than the equivalent Ethernet interface.

As readers of the Banana Skins column in the EMC Journal will know, some manufacturers of products that use HDMI technology have had big problems meeting emissions limits, due to high levels of HDMI emissions spikes. This was eventually found to be entirely caused by poor-quality shield-to-connector termination in the HDMI cables they had purchased to connect to their built-in LCD display, an example of which is shown in Figure 4.

It seems that it is not uncommon for suppliers to provide properly-made full-specification parts when asked for samples, but provide lower-quality parts when they receive an order for production. This can happen even when purchasing well-known UK, European or US branded products, because those suppliers are in turn supplied by less well-known companies (not all of them based in China) and don't do sufficient inspection at their own goods-in.

Just the other week I was told of a military equipment company that needed to buy some obsolete ICs, and – via the Internet – found a supplier in Asia who supplied them with full-spec samples. But when the production quantity was delivered, none of them functioned correctly and X-ray inspection eventually showed that despite appearing to have the correct package and markings, inside they were a variety of different ICs none of which were the correct types.

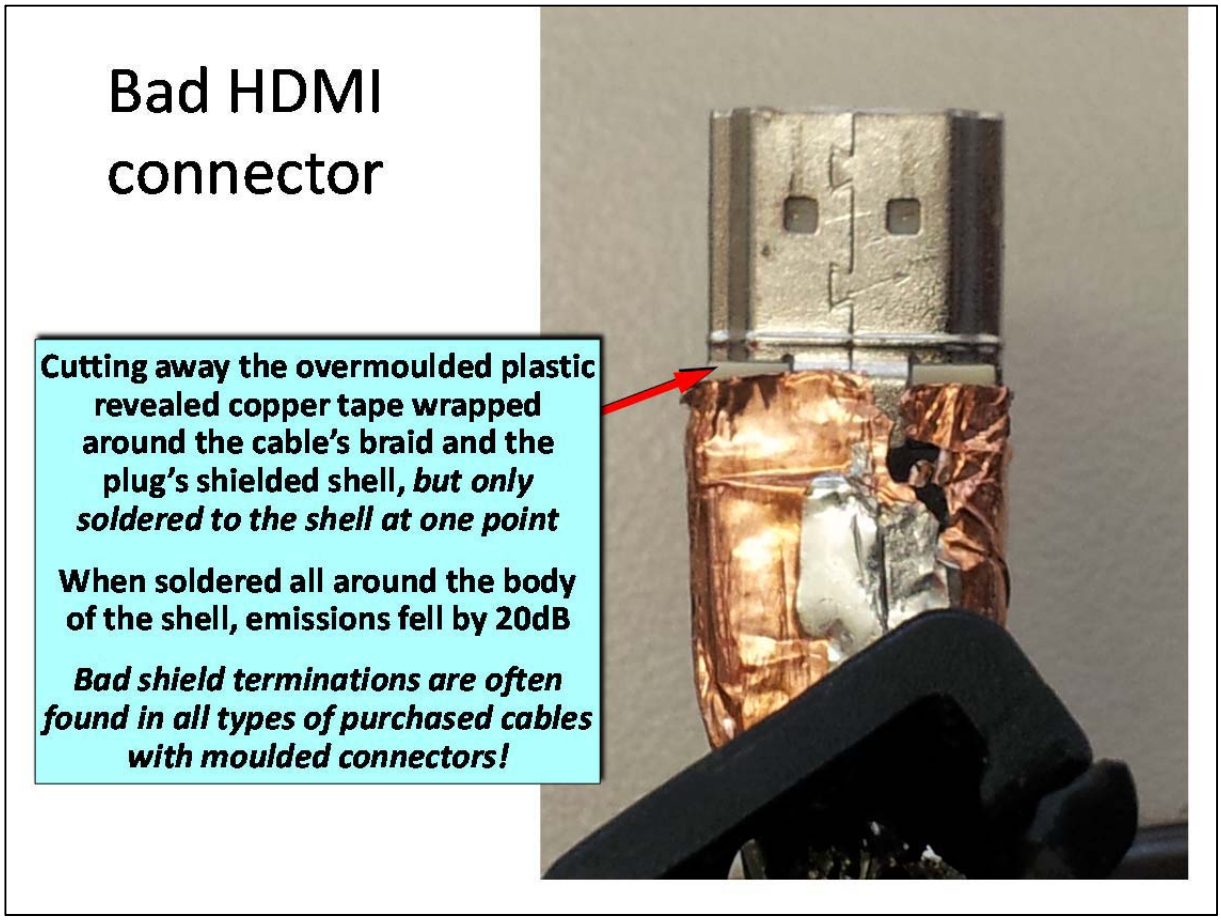

Figure 4 Example of poor shielding in a moulded HDMI cable

It seems that it is not uncommon for suppliers to provide properly-made full-specification parts when asked for samples, but provide lower-quality parts when they receive an order for production. This can happen even when purchasing well-known UK, European or US branded products, because those suppliers are in turn supplied by less well-known companies (not all of them based in China) and don't do sufficient inspection at their own goods-in.

Just the other week I was told of a military equipment company that needed to buy some obsolete ICs, and – via the Internet – found a supplier in Asia who supplied them with full-spec samples. But when the production quantity was delivered, none of them functioned correctly and X-ray inspection eventually showed that despite appearing to have the correct package and markings, inside they were a variety of different ICs none of which were the correct types.

But some industry interface standardisation groups do consider good EMC design. For example, Figure 5, which is taken from [12], shows a proposal for developing Ethernet technology for in-vehicle automotive applications using unshielded twisted-pair wires, and compares the spectrum of Fast Ethernet (i.e. 100BASE-T) – the blue trace – with their new proposal, the magenta trace.

The proposed development described by [12] uses pulse shaping and corresponding receiver equalizer, optimised for typical automotive cable lengths (10 metres), especially to meet the very stringent automotive emission and immunity requirements in the FM radio broadcasting band.

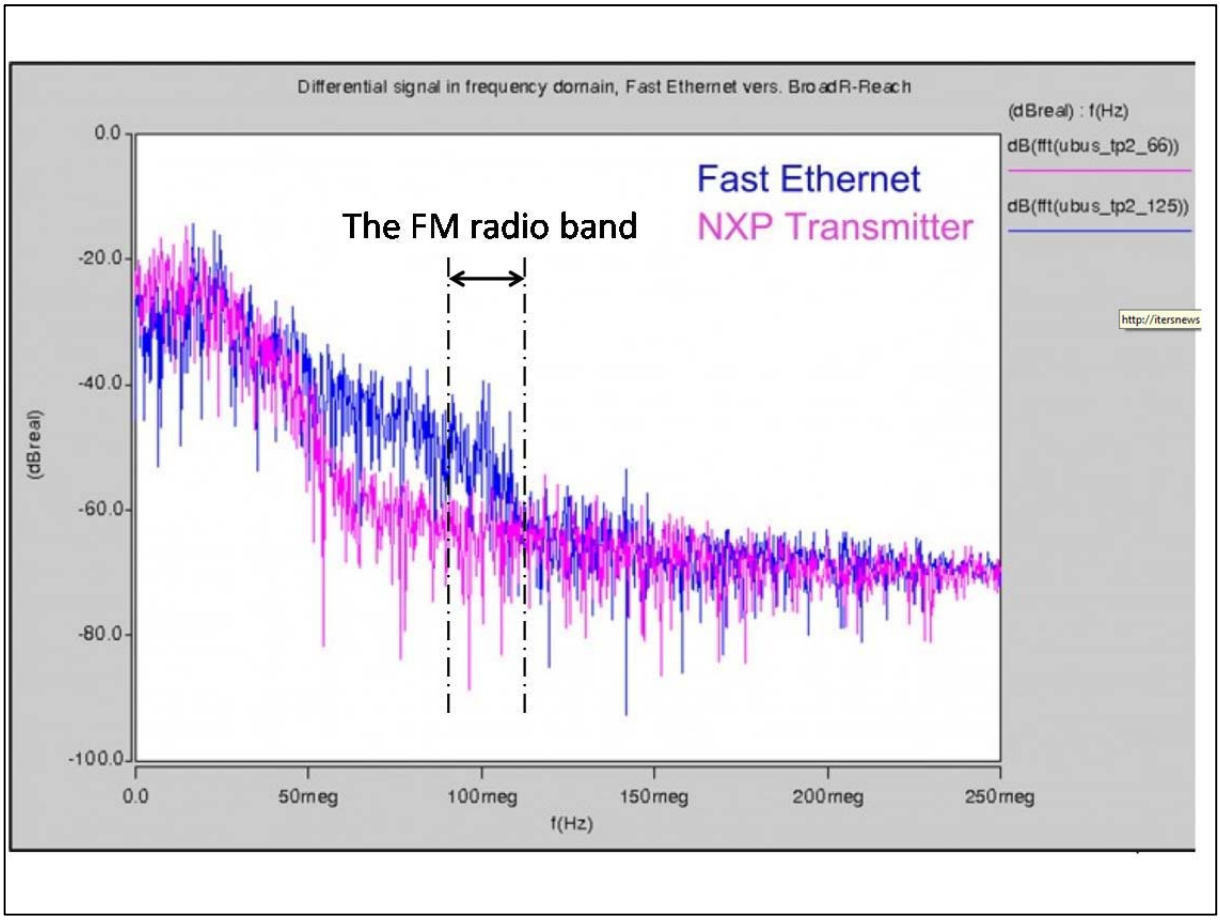

Figure 5 Comparing 100BASE-T signal spectrum with a new 'pulse shaped' design

When either of the emission spectrums from the datacomm designs shown in Figure 5 is measured Quasi-peak, we can expect its level to drop by between 5 and 10dB.

Another example of an industry interface standardisation group considering good EMC design is PCI Express (PCIe), which uses data scrambling to 'whiten' its emissions spectrum to lower its overall cost of implementation.

[13] describes the polynomial algorithm that PCIe uses for data scrambling, which is implemented by means of a 16-bit Linear Feedback Shift Register on the serial NRZ data it sends over each of its many 2.5Gb/s or 5Gb/s 'lanes'.

As well as removing frequency spikes from its signal spectrum, scrambling also helps reduce the significance of crosstalk between PCIe lanes.

3 Spread-spectrum clocking techniques

[21] says the best modulation waveform tends towards a sawtooth (i.e. a discontinuous frequency modulation FM). Optimal discontinuous FM is about 0.8dB better than sawtooth. Triangle wave (continuous) FM is about 4dB worse than sawtooth, and optimised continuous FM is about 2dB worse than sawtooth. Lots of great figures and references in [21].

Spread spectrum clocking does not reduce actual emissions levels at all – only their level when measured using a CISPR 16-1 compliant Measuring Receiver [14].

The measured reduction correlates with a similar reduction in interference with the radio channels that CISPR 16 test methods are intended to protect, but does not necessarily mean that a reduction in interference will occur with other situations.

For example if the occurrence of excessive clock jitter for less than, say, 0.5µs (500ns, e.g. 50 cycles of a 100MHz clock) could cause unacceptable EMI problems in a 'victim' equipment, then none of the CISPR 16 specified detectors will provide a meaningful measurement of the potential EMI.

Spread-spectrum clocking modulates the frequency of digital clock oscillators at a rate of (typically) between 10kHz and 100kHz, so the spikes at the clock's fundamental frequency and harmonics move around rapidly in the spectrum and spend most of their time outside the Measuring Receiver's resolution bandwidth (RBW).

At any given instant, the usual 'spiky' clock harmonic emissions spectrum obtains, but measuring it with CISPR 16's detectors gives a lower measured value because their integrators only charge up while the frequency spikes are within the receiver's narrow bandpass filter (it's RBW) that is centred on the measured frequency. Figure 6 shows an example, taken from [19].

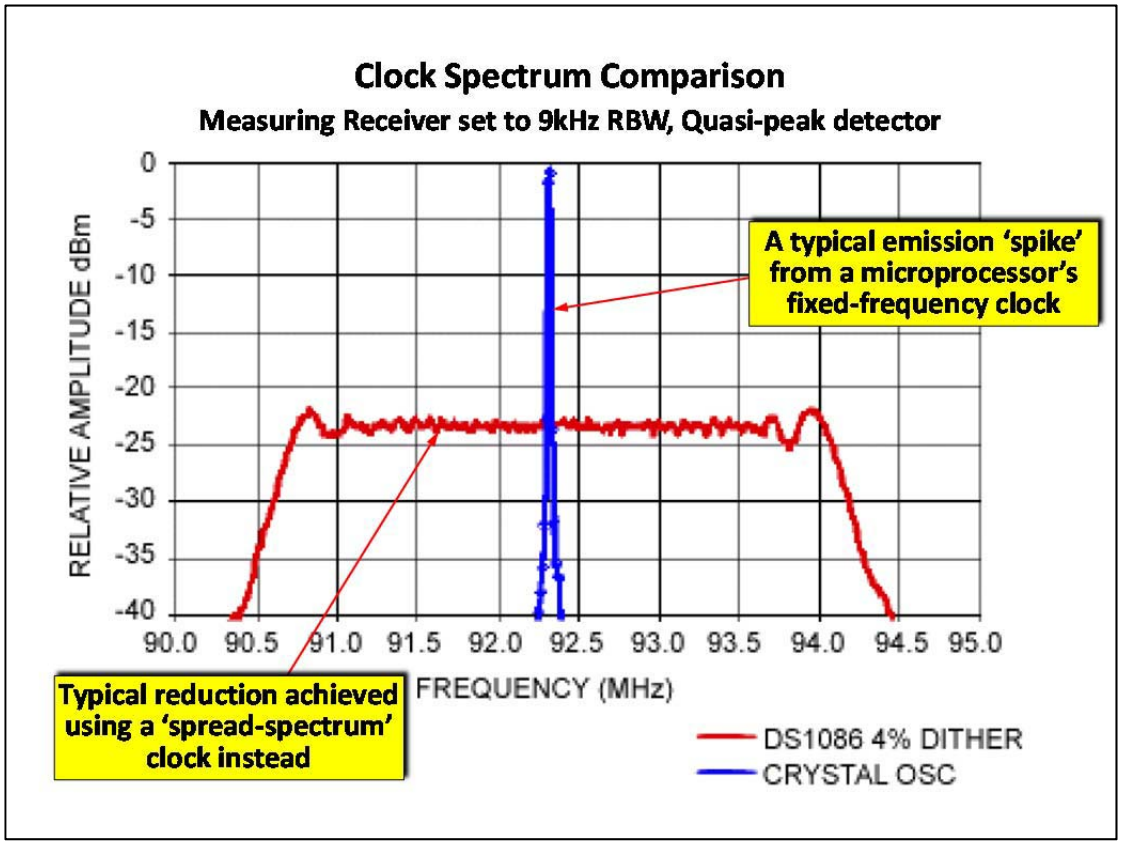

Figure 6 Reducing measured emissions by using spread-spectrum clocking

Some ETSI and ETSI/EN (i.e. telecoms/radiocomms) and some military EMC standards, set limits for emissions as measured with the Peak detector (instead of the Quasi-peak that is normal for IEC EMC standards).

Because even the Peak detectors specified by CISPR 16 use some signal integration, spread-spectrum clocking can still provide some useful reductions in measured emissions, although not as much as can be achieved for the Quasi-peak, Average and RMS detectors, which all use much longer integration times. This is evident from Table 1, which I have extracted from Table 1a in [15].

TABLE 1 Example of measurement results in dB(μV) for the different CISPR 16-specified detectors (Peak, Quasi-Peak, Average and RMS)

Detector type	Measuring the unmodulated RF carrier	Measuring the same RF carrier when it is frequency-modulated (i.e. spread-spectrum)
Peak (PK)	55.6	50.39
Quasi-peak (QPK)	55.4	49.30
Average (AV)	55.38	38.38
RMS	55.38	42.50

When spread-spectrum clocks were first proposed as a way of meeting the FCC's EMC standards whilst spending less on filtering and shielding, almost all radiocommunications, and radio and television broadcasting used fixed-frequency channels.

When the unintentional emission of a clock harmonic, for example from a computing device in the vicinity of a radio receiver happened to coincide with intentional emissions of a radio channel that was to be received, it could easily swamp the receiver's sensitive RF amplifiers and prevent acceptable quality of reception.

Also, for continuous RF immunity testing, spread-spectrum clocking can improve the immunity to certain modes of interference.

When the FCC [16] tested the first computing devices to use spread-spectrum clock oscillators, they found that they did indeed cause permit better quality of radio channel reception, and so the technology was officially accepted. Manufacturers using spread-spectrum clocks could meet the same emissions levels by spending less on EMI mitigation.

However, there is evidence that with modern spread-spectrum/digital/wideband radiocommunications, spread-spectrum clocking of unintentional emitters causes more interference than fixed-frequency clocks, see [15].

It seems that it is best if the interferer's spectrum is spread while the 'victim's' signal uses a fixed frequency, or if the victim's signal spectrum is spread while the interferer's frequency is fixed – but that both should not use either fixed or spread frequencies.

Perhaps, when most radiocomm's and broadcasting uses spread spectrum modulation instead of fixed channels, and CISPR 16 emissions test methods have been changed to properly reflect the effects of interference on them, we will need to change back to using fixed-frequency digital circuit clocking to meet the emissions limits with lower costs!

Note that spread-spectrum clocking cannot be used for every type of signal, or every type of synchronously-switched digital IC.

PCI Express (PCie) can use spread-spectrum clocking for its reference clock, Refclk. 100MHz is a common Refclk frequency, but other frequencies are also used.

The PCIe standard requires a Refclk with better than ±300ppm frequency stability at both the transmitting and receiving devices, and supports three distinct clocking architectures: Common Refclk, Separate Refclk and Data Clocked Refclk.

The Common Refclk architecture has the advantage of permitting spread spectrum clocking, and PCIe devices are specified to reliably transmit data when using this with a spread spectrum modulation rate of between 30 and 33 kHz, and a downspreading ratio of up to +0, -0.5%.

A disadvantage of using spread-spectrum with PCIe is that the Refclk must be distributed to every PCIe device with a clock-to-clock skew of less than 12ns. This can be difficult for large PCBs or when crossing a connector to a different PCB [17].

Figures 7 and 8, which are taken from [18] may be compared to show the reductions in the levels of the frequency 'spikes' at the fundamental and harmonics when a 100MHz fixed-frequency clock is downspread by +0, -0.5% at 32kHz. This would be a typical result when spreading the spectrum of a PCIe Common Refclk running at 100MHz

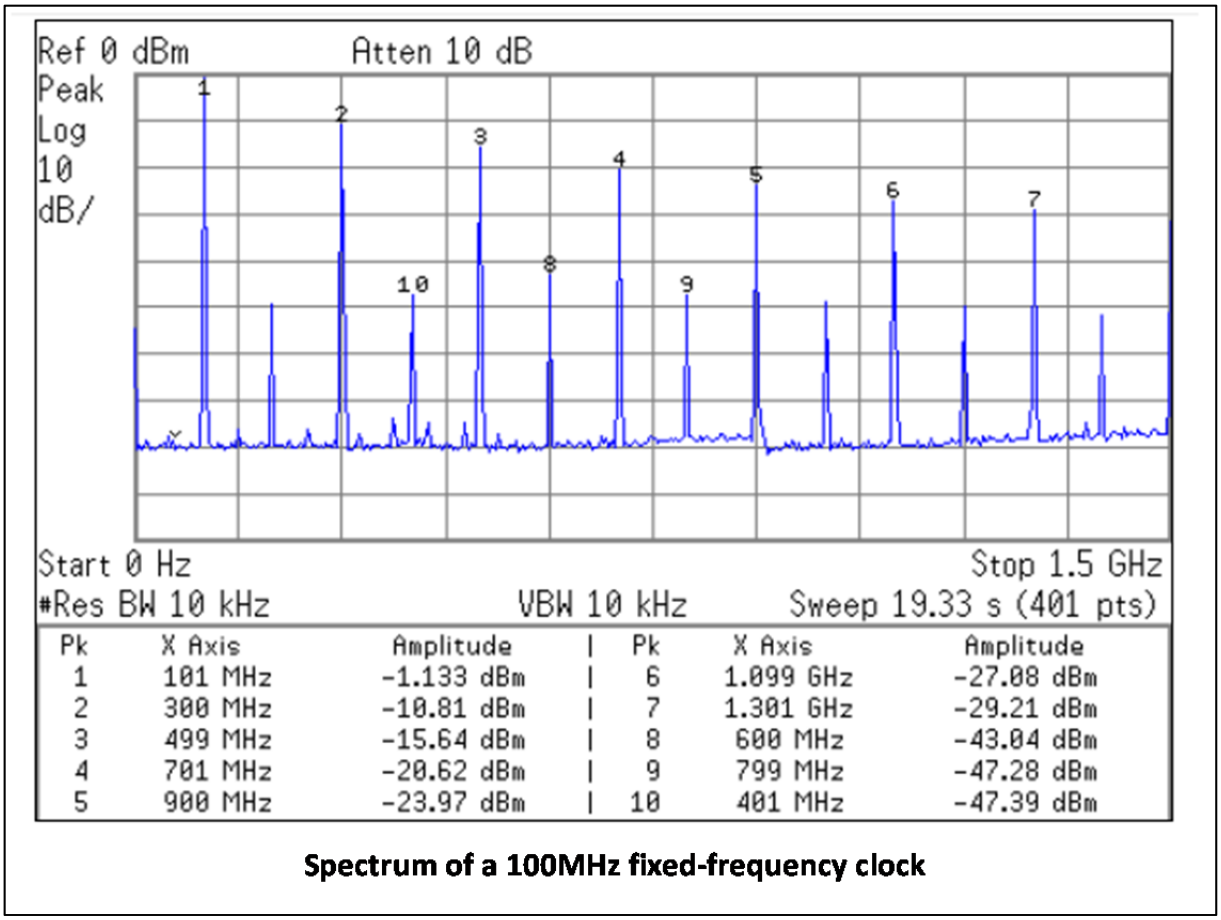

Figure 7 Signal spectrum of a fixed-frequency 100MHz Common Refclk

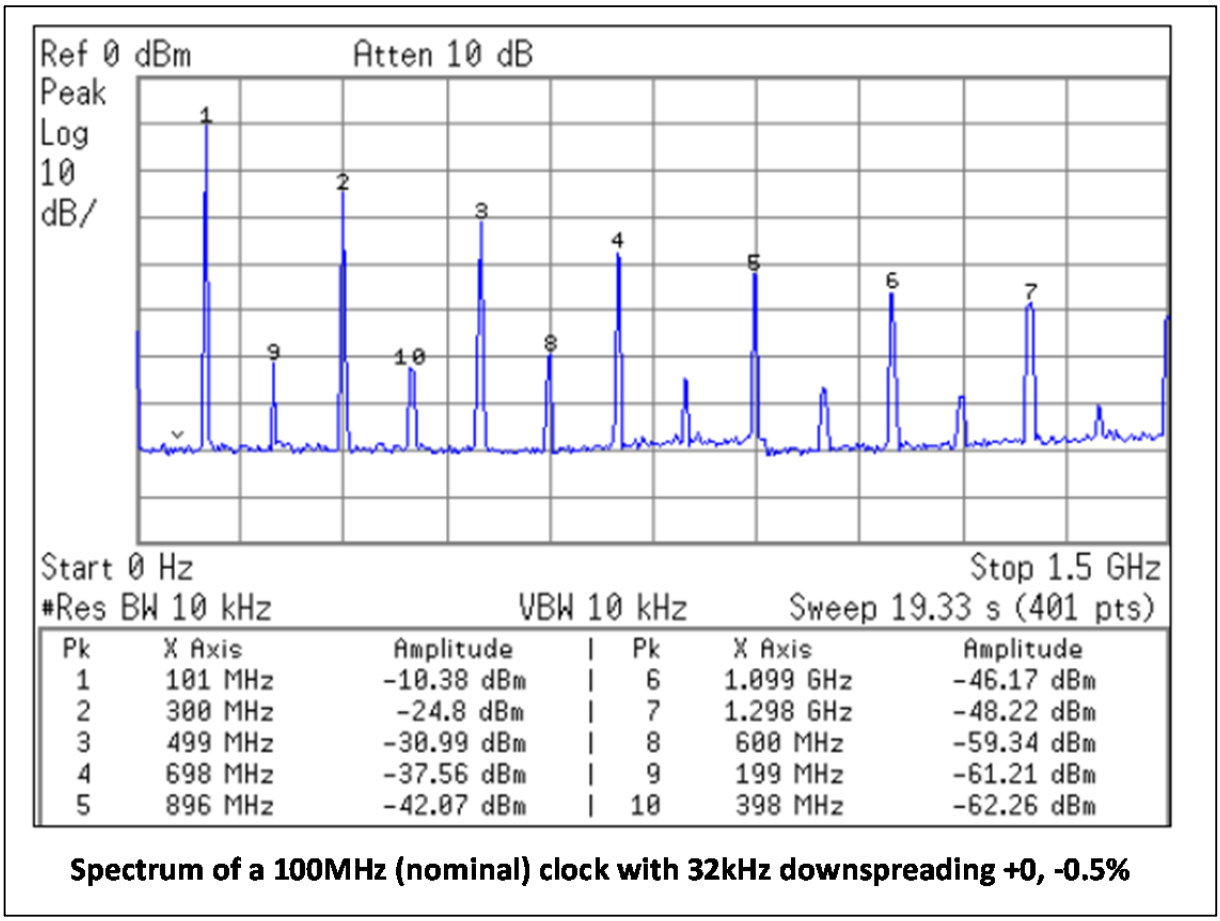

Figure 8 Signal spectrum of a spread-spectrum 100MHz Common Refclk

Spread-spectrum clocking is not restricted to reducing the measured emissions from microprocessors and other digital processing – it is increasingly used in power conversion, where much larger spreading ratios can usually be used than are possible with digital circuits.

A digital clock might be spread (sometimes called 'dithered') by as much as ±2% as shown in Figure 6, but most practical situations use spreading ratios of +0% and around -1%, known as downspreading, because the timing constraints associated with synchronous digital circuits mean that high spreading ratios can cause unreliable functionality.

However, a power converter's switch-rate clock can often be spread by as much as ±10%, maybe more, without causing functional problems. Some converters even use 'chaotic' switching control to try to create more perfectly random emissions than are possible with 'traditional' spread-spectrum clocking techniques.

[19] shows how a DC power converter with an external synchronization pin (e.g., the MAX1703) can allow us to control its frequency with a spread-spectrum clock. Figure 9 shows the conducted emissions spectrum of a free-running power converter, switching at about 300kHz, showing the fundamental switching frequency and its harmonics which extend to 10MHz and more. Peak emissions are -39dBm at its fundamental frequency.

Figure 10 shows the effect of running the exact same converter, under the exact same conditions except that its external synchronisation pin is connected to a spread-spectrum oscillator centred on 300kHz. No discrete (spiky) frequencies are now visible, and the peak is now -53dBm at about 420kHz.

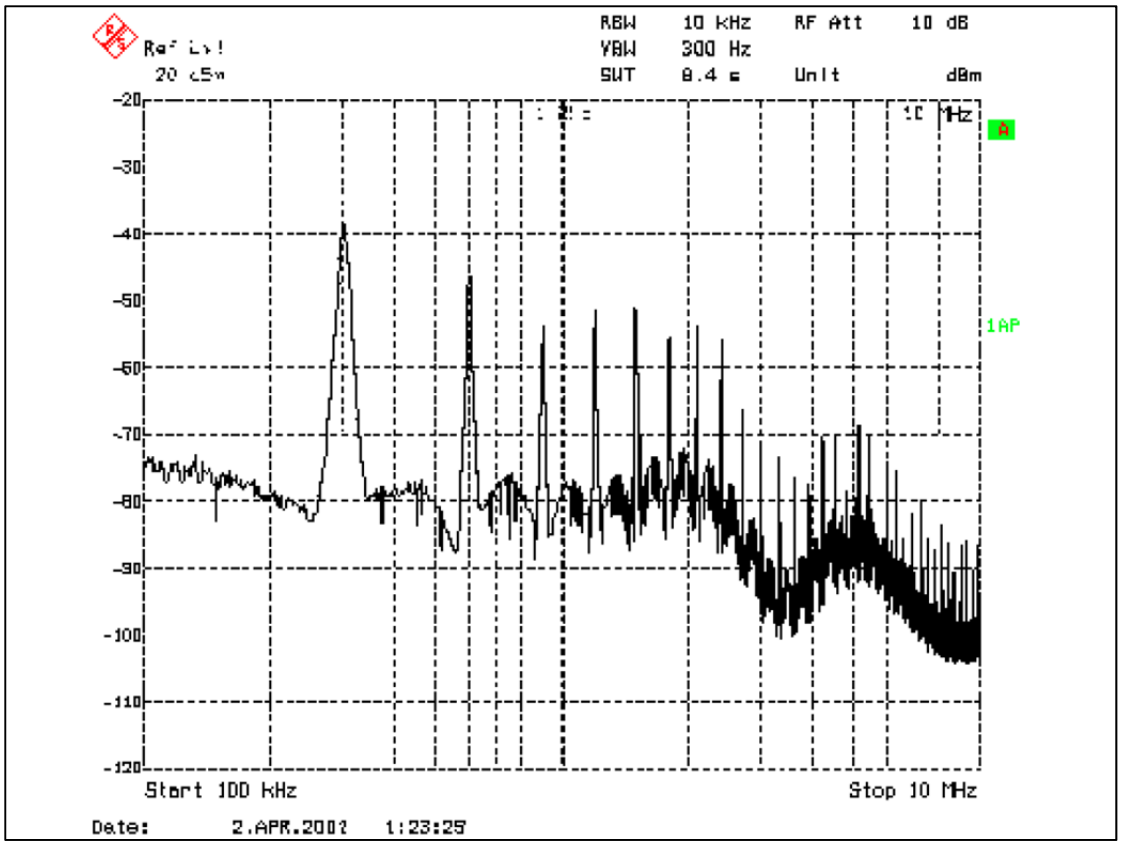

Figure 9 Spectrum of power converter with fixed-frequency clock

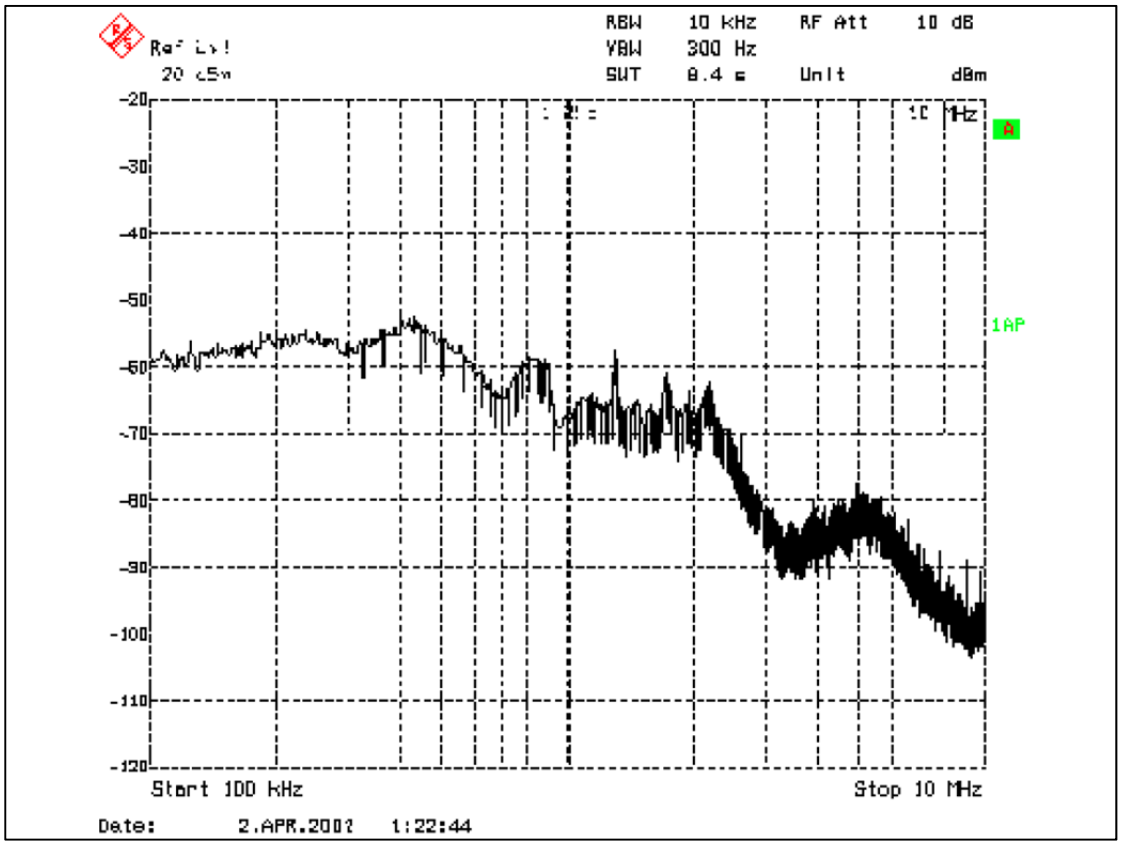

Figure 10 Spectrum of the same power converter with a spread-spectrum clock

It is important to note that the use of spread-spectrum clocking in Figure 10 has caused the 'noise floor' to rise, because the switching energy is no longer concentrated at a few frequency 'spikes', but is spread more uniformly across the measured frequency range.

However, the measurements in Figures 9 and 10 were made with a CISPR 16 Peak detector, and when the measurements are made with Quasi-peak and/or Average detectors we can expect our emissions to be at least another 5dB lower, for example making the 420kHz peak in Figure 7a measure closer to -60dBm.

Figure 11 shows an example, taken from [20], of using a spread-spectrum clock for a PWM-based power inverter. The frequency ranges of the two graphs are both 530 to 800kHz, chosen to make the harmonics of the 17kHz PWM clock more visible.

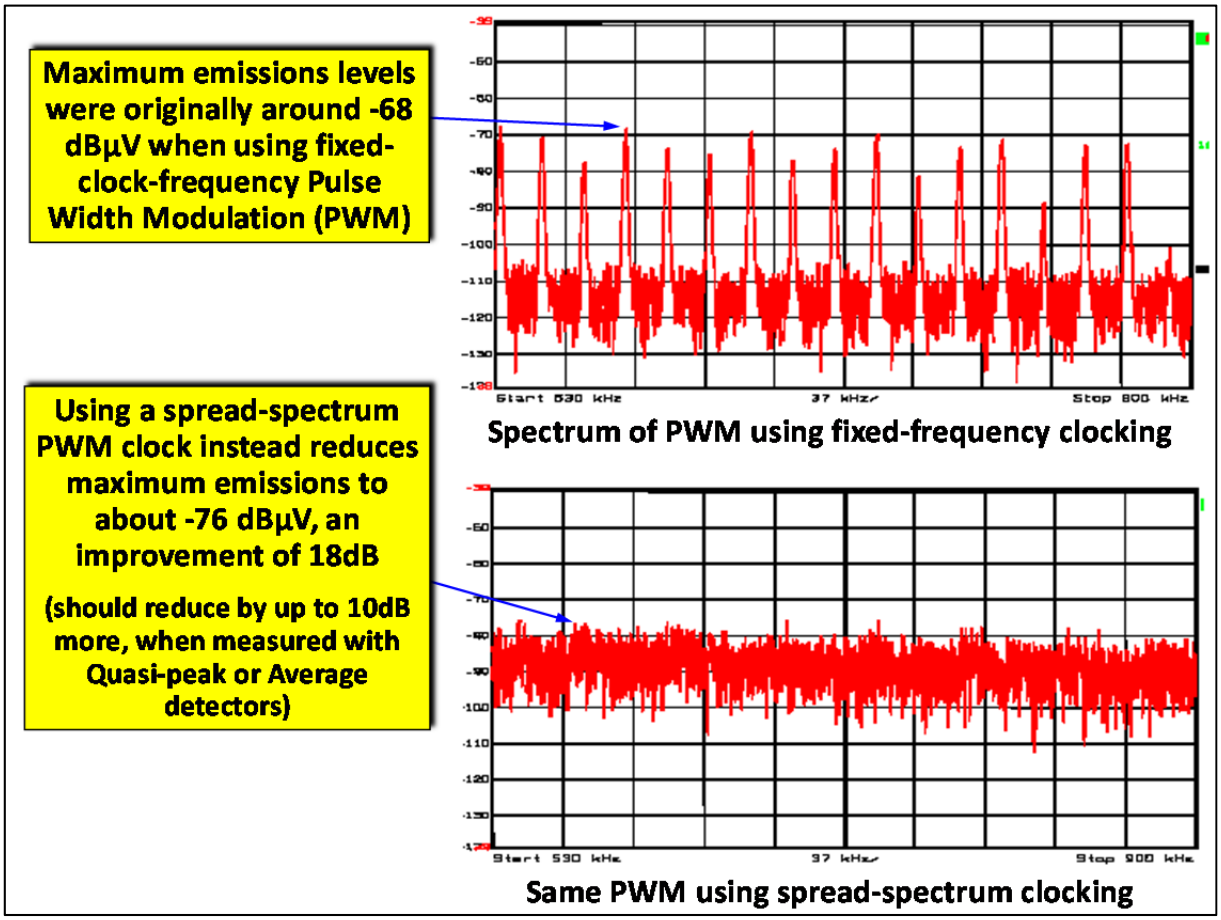

Figure 11 An example of reducing the noise emissions of a power inverter

4 Conclusions

EMC is not just for EMC engineers. Wise Project Managers need to ensure that electronic designers use data scrambling and spread-spectrum clocking (and many other good EMC engineering techniques) where appropriate to help achieve EMC compliance more easily, with lower overall costs, shorter times-to-market, and lower project financial risks.

5 References

[1] “BOM cost and profitability”, Keith Armstrong, The EMC Journal, Issue 82, May 09, from the archives at www.theemcjournal.com.

[2] “The Consumer Electronics Boom”, KPMG, 2008, www.kpmg.com/AU/en/IssuesAndInsights/ArticlesPublications/Documents/The-consumer-electronics-boom.pdf

[3] Keith Armstrong’s many EMC Journal articles on good EMC engineering at all levels in a project, are available from www.compliance-club.com/KeithArmstrong.aspx.

[4] Textbooks on good EMC engineering by Keith Armstrong and Tim Williams are available from www.emcacademy.org/books.asp.
Please note that Keith’s books are only available from this site, and not from Amazon or other sites or shops. Ignore anything on Amazon or anywhere else that says they are out of print – they are colour-printed on demand and so are always available new.

[5] Keith Armstrong’s training courses on proven cost/time-effective good EMC engineering practices for all stages in any type of project, and other EMC or Safety topics, are listed at www.cherryclough.com. They have always received appraisal scores of at least 80% satisfaction from their attendees, all over the world.

[6] “Gigabit Ethernet 1000BASE-T - Focus - Physical Coding Sublayer (PCS), Functional Basics and Overview, Based on: IEEE 802.3ab”, Bob Noseworthy, 1000BASE-T Tutorial Series, IOL Gigabit Ethernet Consortium, 19 November 1998, updated 29 March 2000, https://www.iol.unh.edu/services/testing/ge/knowledgebase/pcs.pdf

[7] “802.3ab – A Tutorial Presentation”
http://grouper.ieee.org/groups/802/3/tutorial/march98/mick_170398.pdf

[8] “An Introduction To TMDS – Understanding DVI-D, HDMI and DisplayPort Signals”, https://www.bicsi.org/pdf/Regions/charlotte_se_4_11/Understanding%20DVI-D,%20HDMI%20and%20Display%20Port%20Signals.pdf

[9] “Gigabit Ethernet over Category 5”, Fanny Mlinarsky, Scope Communications, Inc., 6 Jun 1998, www.octoscope.com/English/Collaterals/Articles/octoScope_GbitEthernet_GEC98_19980606.pdf

[10] http://en.wikipedia.org/wiki/Non-return-to-zero

[11] “Effects of Ferrite Core on EMI and Signal Transmission for HDMI”, Yuang-Shung Lee and Hai-Hwa Hu, www.egerate.com/articles/Ferrite_Core_HDMI.pdf

[12] “Automotive Ethernet, a holistic approach for a next-generation in-vehicle networking standard”, Peter Hank, Thomas Suermann and Steffen Müller of NXP Semiconductors Germany, August 9, 2012, http://itersnews.com/?p=10541

[13] “PCI Express System Architecture”, by Ravi Budruk, Don Anderson and Tom Shanley, Addison-Wesley Professional, 2004, ISBN-10: 0321156307, ISBN-13: 9780321156303

[14] CISPR 16-1, “Specification for radio disturbance and immunity measurement apparatus and methods - Part 1: Radio disturbance and immunity measuring apparatus”, from http://webstore.iec.ch/webstore/webstore.nsf/artnum/031863.
A useful article is “The EMI Receiver according to CISPR 16-1-1” by Dipl. -Ing. Dieter Schwarzbeck, from www.schwarzbeck.de/appnotes/EMIRcvrCISPR16.pdf

[15] “Possible Effects of Spread-Spectrum-Clock Interference on Wideband Radiocommunication Services”, Manfred Stecher, IEEE, http://ieeexplore.ieee.org/stamp/stamp.jsp?arnumber=01513472

[16] The Federal Communications Commission in the USA, which first introduced regulations on emissions in 1983, www.fcc.gov

[17] Silicon Laboratories, Inc. Rev 1.0 1, “Selecting the Optimum PCI Express Clock Source”, www.silabs.com/Support%20Documents/TechnicalDocs/PCIe-Clock-Source-Selection.pdf

[18] “PCI Express – Signal Integrity and EMI” ANTC203, Micrel Inc., www.micrel.com/_PDF/HBW/App-Notes/ANTC203.pdf

[19] Maxim-Dallas App Note 3512: “Spread-Spectrum Clock Oscillators Lower EMI”, 29 April 2005, http://www.eetasia.com/ARTICLES/2005JUN/A/2005JUN08_ACC_POW_AN.PDF?SOURCES=DOWNLOAD

[20] “Reduction in Radio Interference Through PWM Frequency Dithering”, L Comstock, H Rakouth, J Ruiz, IEEE International EMC Symposium, Portland, Aug 14-18 2006, ISBN: 1-4244-0294-8

[21] “Optimal Discontinuous Frequency Modulation for Spread-spectrum Clocking, David De Caro, IEEE Trans. EMC, Vol 55, No 5, October 2013, pp 891-900 (in turn contains a great many relevant references)