

The water-bed and the leaky bucket

Tim Williams
Elmac Services
Wareham, UK
timw@elmac.co.uk

Abstract—The common situation of EMC mitigation measures having the opposite effect from what was intended, is described, with particular regard to clock harmonic radiated emissions. Two mechanisms for the contradictory effects are offered: changes in the harmonic structure of the source circuit or device, and phase cancellation of fields from multiple source structures in the product. Measurements of a simple emitting device are seen to be reflected in the Fourier transform of circuit waveforms and the modelled field patterns of its equivalent antenna structure.

I. INTRODUCTION

One of the most commonly-encountered phenomena when a new product is being tested for radiated emissions is what has become known as the “water-bed” effect. The most typical illustration of this effect is when a set of harmonics from a particular clock source is being measured and some of the harmonic emissions are over the relevant limit. Various mitigation techniques – filtering, shielding, schematic or ground structure modification – are tried to reduce this excess, but in every case the modification drops one emission frequency only to cause a different harmonic to pop up over the limit. It is just as if the harmonic structure were sitting on a water-bed: pushing one area down only results in another area increasing. Often, when an azimuth scan is repeated, the emission levels have dropped in one direction but increased in another. What is going on, and what can be done about it?

II. HARMONIC STRUCTURE OF SOURCE DEVICE

Two mechanisms can be invoked to explain the effect. The first relates to the source in the circuit, without considering its radiation. The harmonic structure of a single-frequency clock depends on the detail of its waveform, particularly rise and fall times, which in turn will vary with the high frequency impedance of its load circuits and associated parasitics, as well as variations in supply voltage and operating temperature. Changing any of these will affect the relative amplitudes of the harmonics; some will decrease, but some will increase. For instance, loading a clock driver with series impedance will normally reduce the amplitude of the higher order harmonics but could, through the change in load, increase that of low orders. What is more, the coupling may not in fact be from the signal circuit but from associated ground or power rail currents. In this case, circuit changes can not only affect the structure of the power current harmonics but also, for instance through decoupling placement, re-route their current paths.

A. Investigation of a simple emitter

To investigate the effect a circuit was constructed consisting of a 40MHz clock oscillator driving a 74HC244 tri-state buffer device. The buffer could be connected directly to a spectrum analyser, or to a length of ribbon cable; its outputs could be selected high impedance, on but static, or on and driven with the clock signal. The clock source was permanently active and driving the buffer inputs. The schematic is shown in Figure 1. The PCB layout for the circuit was deliberately designed to be poor from the RF point of view so that the 40MHz clock would be radiated through both the circuit and the cable. Figure 2 shows the spectrum profile for the harmonics when the analyser is connected to the buffer output through a resistive attenuator.




Figure 1 Simplified schematic of the test item

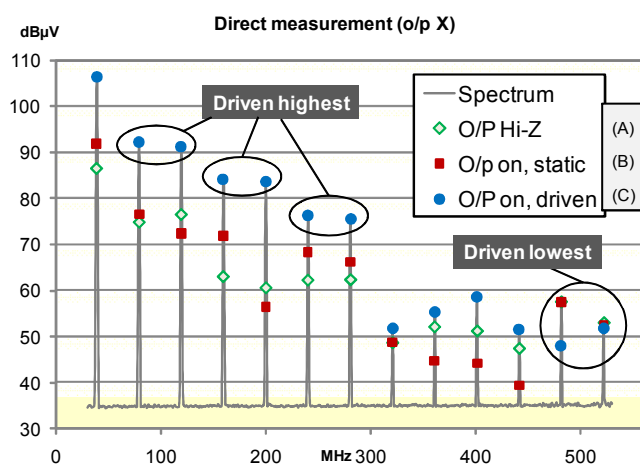


Figure 2 Changes in harmonic structure of the 74HC244

The three cases are (A) the device has both sets of output disabled, (B) output 1 is driven on but permanently high, (C)

output 2 is driven from the clock. In neither case (A) nor (B) is the clock intended to appear directly at the analyser, but the measurement is of the stray levels which are developed between the device's output pin and the 0V rail. These levels reflect the poor quality of the PCB layout as well as the impedance of the package connections – essentially it is a measure of the “ground bounce” or simultaneous switching noise of the circuit [1][2], developed partly across the internal impedance of the IC leads and partly across the PCB tracks. (While it is to be hoped that real PCB layouts are not as bad as this, experience shows that they still exist.)

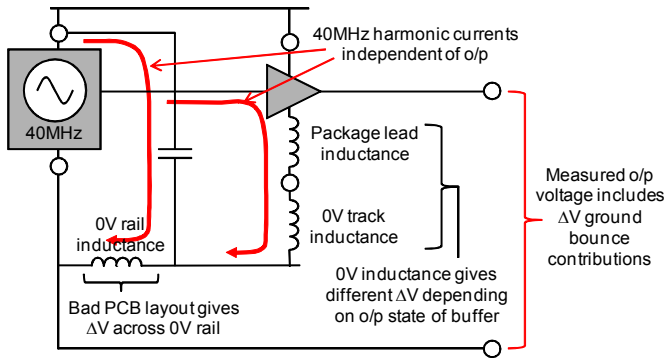


Figure 3 Ground bounce

From this measurement it can be seen that changes due to different drive conditions are contradictory: although in most cases the output is highest when the clock is actively driven, when the output is in tri-state or driven high there are inconsistencies. For the lower order harmonics below 300MHz the changes are as expected. The level is highest when the clock is driven to the output, and the tri-state condition generally has the lowest levels. But at 480MHz and 520MHz the lowest level is when the output is driven on with the clock signal.

The measurement is made through an attenuator of 10dB which is connected between the output pin of the 74HC244 and the 0V rail. But because of the deliberately poor PCB layout, and the fact that the device is in a DIL package and socketed, there is an excess inductance of around 10nH which appears as a common impedance between the device's 0V terminal and the 0V of the rest of the circuit. This passes some internal clock current even when the device is not driving the clock to the output, since the clock signal is present at the input of one section of the device all the time. In addition, poor PCB layout allows 40MHz currents created elsewhere in the circuit, such as the clock oscillator decoupling, to create ground voltage drops which are added to the measured output (Figure 3).

The voltage developed across these inductances (the ground-bounce voltage) is passed to the output, and the waveforms are modified by the internal state (tri-state or driven) of the device [3].

B. Fourier analysis of the driving waveform

A Fourier analysis demonstrates that even quite small variations in the likely ground-bounce waveform can give the kind of effect noted in Figure 2. The waveforms of the three output states are shown in Figure 4, and the Fast Fourier Transform of these waveforms is shown in Figure 5. Compare these points with Figure 2, made on the same circuit node but

with a spectrum analyser; the first few harmonic numbers are essentially the same, but higher orders show diverging results. There is a point at which the output driven signal level is less than the output off level, but in this case it is at the 11th harmonic, 440MHz. The 12th and 13th harmonics are quite different from those seen on the analyser. This illustrates the sensitivity of the higher order harmonic amplitudes to small variations in the captured waveforms.

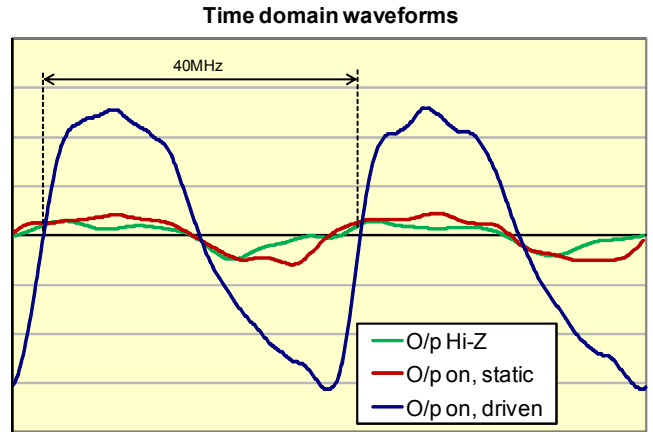


Figure 4 40MHz waveforms at output X




Figure 5 FFT of the 40MHz waveforms

If this one signal drives a dominant emitting mechanism in common mode such as a connected cable or a chassis structure, whose resonances perhaps enhance the emissions at some frequencies, then decoupling or filtering modifications which change the ground-bounce waveform – but don't necessarily attenuate it – will in turn create variations in the radiated profile. As seen above, these variations may contradict the intended and expected improvement in the emissions level, because of phase cancellation effects in the harmonic structure.

III. PHASOR ADDITION FROM MULTIPLE SOURCE STRUCTURES

A second factor to appreciate is that a given emitting source – say, a particular system clock and its harmonics – is almost never radiating from just one point. Instead it generates a driving signal which is distributed across one or more PCBs in

the product, and which can also drive either a differential or a common mode current into one or more connected cables. Depending on the frequency and the dimensions of the PCB(s) and cables, the dominant emitting structure for different harmonic components may vary; and there may in fact be no absolutely dominant structure, with several areas of the product contributing more or less equally to the far field emissions profile. It is for this reason that near field probe checks cannot properly represent far field measurements.

In this situation, the relative phases of the fields emitted by the contributing structures can become important. If one contribution dominates – say, is more than 6-10dB over all others – then phase variations between the separate contributors will be largely irrelevant. Whatever their value and however they change, they will make only a few dB difference to the outcome. But with equal radiating efficiency from multiple contributors, significant notches in the emissions profile can occur when their phases cancel. (This property is of course the basis for the design of phased-array antennas [4], but it is not often appreciated by product designers in EMC work.)

The test object described earlier can also be used to demonstrate this effect. A measurement of radiated emissions at 3m distance over the frequency range from 150 to 650MHz is shown in Figure 6. Different markers show the emissions with the output driven with a clock signal, and with it not driven but disabled to a high impedance. Filled markers show the values when a 1m length ribbon cable is plugged in to the appropriate connector, open markers show no cable. The setup was horizontally polarized, with the cable and unit in the same plane as the measuring antenna and broadside on to it, which configuration is expected to produce the maximum emissions.

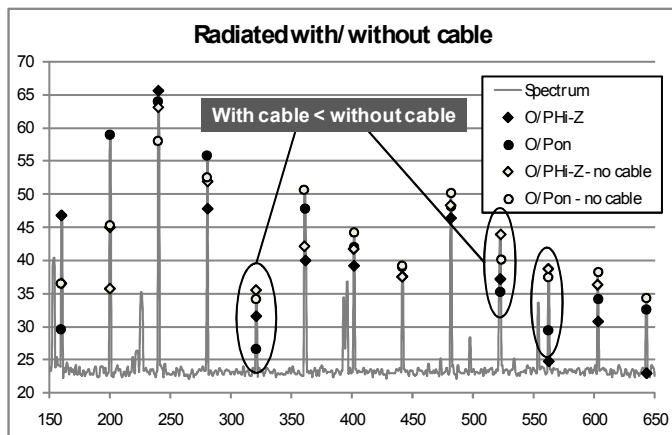



Figure 6 Radiated emissions: with and without cable

Figure 7 Radiated emissions with cable, with/without ferrite



Conventional wisdom would expect that when the cable is connected the emissions would go up, and this indeed happens at the lower frequencies, below 250MHz. But at other frequencies the reverse occurs; this is especially obvious at 320MHz and 560MHz, but is also evident at other frequencies to a lesser extent. At these frequencies, if the cable emissions were to be attenuated by adding a ferrite sleeve as would be typical advice, the total emissions would go up, not down. To verify this, instead of removing the cable, it was partly decoupled with a pair of ferrites in series at the connector end (Steward part no 28R1101-000). The result is shown in Figure 7.

This shows a very similar, but of course not identical, result to the removing of the cable in Figure 3. Again, the effects at 320MHz and 560MHz are reversed from what would be expected. At most other frequencies above 300MHz the levels are equivalent to the removal of the cable. At 160 and 200MHz the ferrite is less effective, suggesting that the cable impedance is rising at these frequencies by comparison with the ferrite impedance – as might be expected for a half-wave open ended cable.

A. Modelling the effect

The complete assembly of PCB and cable is simple enough for it to be possible to model its main features and demonstrate a similar effect through the model. The PCB is single-sided with thin tracks, which not only makes it highly emissive but also allows it to be represented by a wire structure, along with its connected cable. To create the model, the circuit schematic of Figure 1 is reduced to two sources; one represents the 40MHz oscillator driving a signal around a loop on the PCB to the input of the 74HC244 buffer, which is assumed to be purely capacitive. The other represents the output of the buffer driving the connector pins, to which the cable may be connected or not. The two sources will of course have the same frequency and phase relationship, and for the purposes of the model they can be regarded as having the same amplitudes, since they are both CMOS output level devices. A simple diagram of the resulting model is shown in Figure 8, from which it becomes clear that the total radiated field will be due to the combination of three components:

- the small loop driven by the oscillator,
- the long cable driven differentially by the buffer output, and
- the common mode excitation of the cable-and-loop structure by the voltage developed across the loop (the ground bounce potential).




Figure 8 Structural model of the circuit in Figure 1

The wire structure with its sources and loads can then be

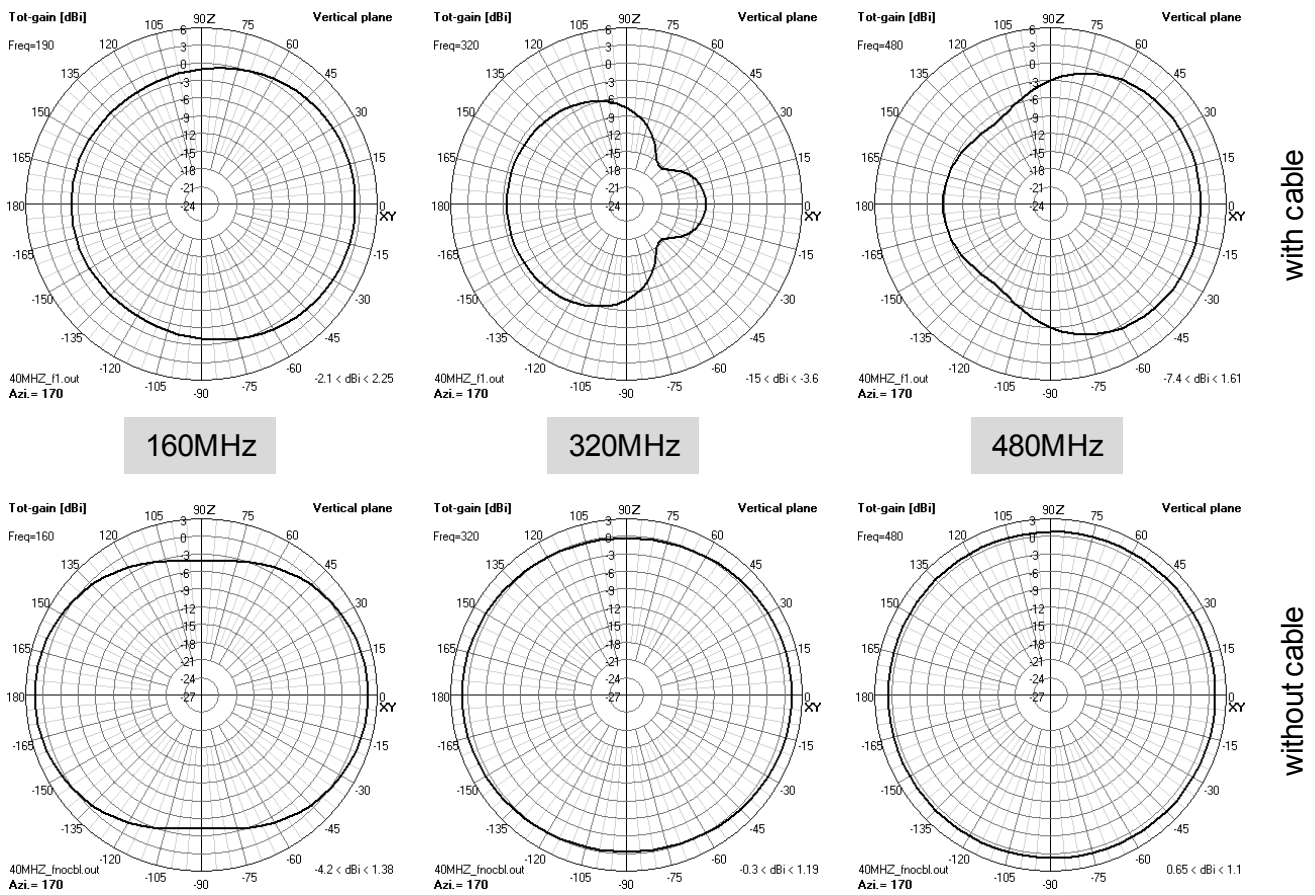


Figure 9 Calculated vertical plane far field patterns with and without cable, at three frequencies

used as the input to an antenna modelling code such as NEC [5]. One attraction of this approach is that NEC will compute the currents on each wire segment as a result both of the driving sources and the mutual coupling between segments, and so takes into account the changes in impedance caused by this mutual coupling. For the purposes of this discussion, we are not interested in absolute values of radiated field but in the differences in the structure's radiating efficiency over frequency, and in how this changes when the structure is varied, in particular by adding or removing the cable.

The far-field pattern results are shown in Figure 9. This gives the calculated far field distribution at three spot frequencies, in a vertical plane which corresponds to the direction of the receiving antenna in the measurement. The three frequencies are 160MHz, 320MHz and 480MHz and the patterns are shown for two cases, with the wire structure corresponding to the cable present or absent.

The most noticeable feature is the striking difference between the plots with and without the cable at 320MHz. The far field pattern is reduced when the cable is attached by around 10dB in the direction of the measurement, whereas this does not happen at the other frequencies. Without a cable, the pattern is largely non-directional, as one would expect from a small loop in the geometry under consideration, in which the loop is the dominant radiating source. But adding in the second source attached to a long cable, the phasing effects and consequent directivity and reduction in emissions in the

direction of measurement become significant at certain frequencies.

This phenomenon will also explain why a reduction in emissions at a particular frequency in one direction can be negated by a corresponding increase in another direction. A change in the current flows in the total radiating structure simply changes the directional response of the structure, without reducing the overall radiated energy.

IV. DISCUSSION: THE LEAKY BUCKET

The above two mechanisms (harmonic phase effects in the Fourier spectrum, and antenna pattern phase effects at different frequencies) have been discussed as if they are separate and unrelated. In fact, they act together, and what is more, one affects the other. Varying the antenna structure changes the output loading of the sources, which in turn changes the distribution of harmonics in the Fourier spectrum, particularly for the higher orders. For instance, the NEC model results don't predict the abnormally low measured level at 160MHz in Figures 6 and 7 when the output is driven on. Although the modelling shows the effects of structural changes on the antenna pattern for a given fixed amplitude of the harmonic spectrum and no phase differences between the harmonics, to properly model the emissions levels would demand that the harmonic amplitudes and phases were individually recreated in the NEC model for each frequency, and iterated until the loading effects were correctly replicated. While this is possible,

the effort involved even for a simple device such as the test object described here would be excessive.

A consequence of this is that in EMC mitigation it is never adequate to say, as is often the temptation, that “we tried that and it didn’t work”, with the implication that that particular fix needn’t be tried again. It is always necessary to have in mind the physics of a mitigation method; if it didn’t work in one set of circumstances, that is useful diagnostic information; and if the physical basis is sound, it may well work after other changes have been applied.

A good analogy for emissions mitigation in general might be an elderly water bucket with several holes of different sizes [6]. The water will pour out through *all* the holes. If you stop up a small hole and leave the big ones, you won’t notice a difference in the leak rate; but once you have fixed the big holes, it will be worth tackling the small ones. But with radiated emissions there is a catch. Because of the phasing effects discussed above, if two holes are of similar size, it’s possible for you to stop one of them and yet *increase* the leaks from the whole bucket. Try explaining that to a country farmer! (You could say that one hole is leaking back into another, but that’s stretching the analogy a bit far).

V. CONCLUSIONS

The purpose of this discussion has been to show that the “water-bed effect” has a predictable, if not entirely simple, foundation in both circuit and electromagnetic physics. Product designers are confused and disappointed when universally-recommended modifications to mitigate emissions are found to

have the opposite effect. An understanding of the general mechanism involved in creating the signals will show why this can happen: for clock emissions, the mechanism can be due to phasing effects both within the harmonic structure of the driving source and related to summation and cancellation of emissions from different radiating structures, driven by the same source.

Without considerable effort to accurately represent all the contributing factors, a model won’t be able to predict the outcome of individual mitigation solutions. It is always necessary to implement good practice throughout a design to keep the amplitudes of all driving sources to a minimum, and particularly to minimize the levels of ground bounce noise, which once created is the hardest to control.

REFERENCES

- [1] R. Senthinathan and J.L. Prince, Simultaneous switching noise of CMOS devices and systems, Kluwer Academic Publishers, 1994
- [2] P. Heydari, M. Pedram, "Ground bounce in digital VLSI circuits", IEEE Transactions on Very Large Scale Integration (VLSI) Systems, Volume 11 Issue 2 pp 180-193, July 2003
- [3] A. Vaidyanath, B. Thoroddsen, J.L. Prince, "Effect of CMOS Driver Loading Conditions on Simultaneous Switching Noise," IEEE Trans. on Components, Packaging, and Manufacturing Technology-Part B, vol. 17, no. 4, pp.480- 485, Nov. 1994.
- [4] C. A. Balanis (ed), Modern Antenna Handbook, Wiley 2008
- [5] 4NEC2, Arie Voors, <http://home.ict.nl/~arivoors>
- [6] This analogy was first proposed by Ulf Nilsson of EMC Services, Sweden, in 1990