IMMUNITY FROM AIRBORNE AGENTS

Testing the interference performance of radio data links.

Nick Long
Great Circle Design

Abstract

Any radio system can be disrupted by interference. The end user naturally expects a reasonable degree of interference immunity in the design. In many cases, though, there is no definition of what is reasonable and no agreement on test methods.

The interference performance of radio equipment sold to non expert customers can range from downright negligent to surprisingly good.

This paper reports on development of standardised and automated test methods for short range telemetry links. Results obtained from several commercially available devices are presented, showing a wide variation in the ability to reject interference.

Introduction

A description of the performance of a radio system can run to many pages of information, whether it is a specification or a list of measurements. While many of the parameters relate to the normal operation in benign conditions, a surprisingly large proportion are in the list only because they matter when conditions are difficult, ie., when there is unwanted interference.

Receiver parameters such as bandwidth, co-channel rejection, adjacent channel rejection, intermodulation performance, intercept points, blocking level and spurious response are routinely banded about, but they only reason they are relevant is that they give us clues as to the behaviour in the presence of interference.

While it easy enough to write these parameters into a specification, it is somewhat more difficult to measure them in practice, and even then they do not give us a particularly good picture of the interference immunity of the receiver.

This paper describes a way of bringing all these parameters together and showing the interference immunity of a receiver in one simple manner. The necessary measurements can be made manually, but they are well suited to being automated.

The Well diagram

The well diagram is so called because of its shape. It is the interference level that a receiver can tolerate plotted against frequency. The diagram for an ideal receiver would have the shape in Figure 1.

Note that the depth of the well is not the receiver sensitivity. It is the limit of co-channel interference for the particular wanted signal level chosen for the test. Much more important than the depth of the well is how far the rest of the diagram differs from the ideal. It is necessary to define the level of wanted signal for the test, as both the inner and outer parts of the diagram are affected by it.

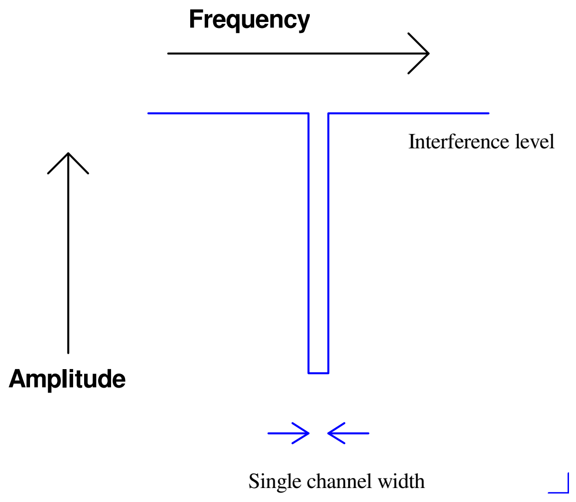

Figure 1

Measuring any one point on the plot is relatively simple; with a basic AM or FM receiver it can be done with two signal generators. There are, however, some important considerations.

Isolation and screening. With higher performing receivers, we can be dealing with a combination of very low level and very high level signals. Attention to system layout and screening is important to ensure the levels really are what we think.

Pass/Fail Criterion. Defining the point of failure as the interference rises can be difficult, particularly if we have only the receiver and not the full system. In a complete data link we can look for packet errors or loss of clock recovery, but examining an IF or baseband signal on a scope is subjective. Fortunately, in many cases the change from perfect to useless is rapid and this subjectivity leads to only a few dB spread.

Order of signals. The simplest test is to apply the wanted signal and then raise the unwanted level until a problem is notes. This can often be misleading; in reality the unwanted signal may already be there before the wanted one. This can be very important in the case of receivers using PLL demodulators or similar systems that lock on to a signal. I have seen an extreme case in a low cost receiver where the PLL was used not only for FM demodulation but also for the final IF filtering. Wanted signal first testing showed narrow bandwidth and good adjacent channel rejection, but unwanted signal first revealed appalling susceptibility over many MHz. Even without these problems, issues such as clock recovery and AFC still make unwanted signal first testing the most representative.

A good procedure is this:

Apply unwanted signal
Perform “ping” test with wanted signal
Remove wanted signal
Raise level of unwanted signal one step and repeat.
Remove unwanted signal between frequency changes.

An even better one is to test with both rising and falling unwanted levels. Unfortunately, these procedures make testing slower and more tedious.

What follows is some results from measurements on data systems for the 433 MHz band. This is a licence exempt band with minimal regulation; the radio regulations cover only the transmitters and there is no specification on receivers other than what, if anything, may be agreed between the supplier and the customer. Virtually all the equipment here can be described as "low cost", but there is a huge variation in the performance.

Also as we go through these, it will be seen how the measurements evolved from ad hoc manual to fully defined and automated.

Example 1

The first example is a very low cost AM receiver intended for general purpose low data rate applications. Figure 2 shows the well diagram. This looks reasonably good until we realise that the horizontal axis is 10 MHz/div. Figure 3 shows it on a more representative scale.

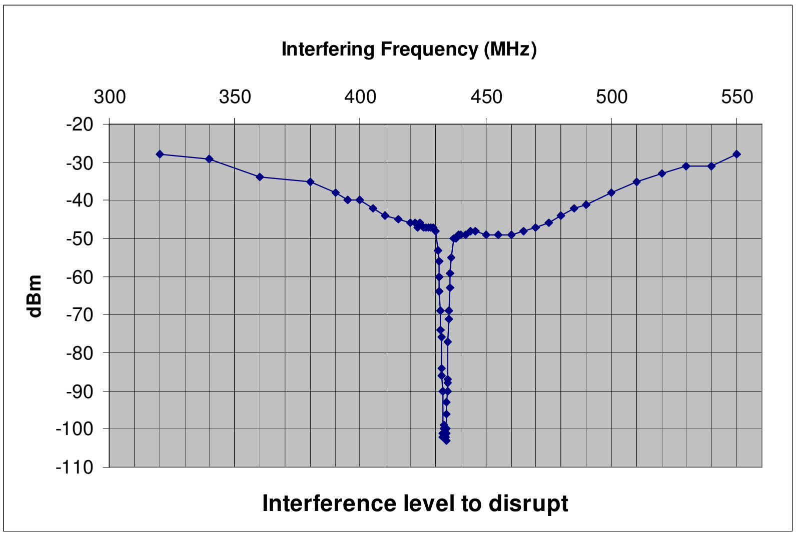

Figure 2

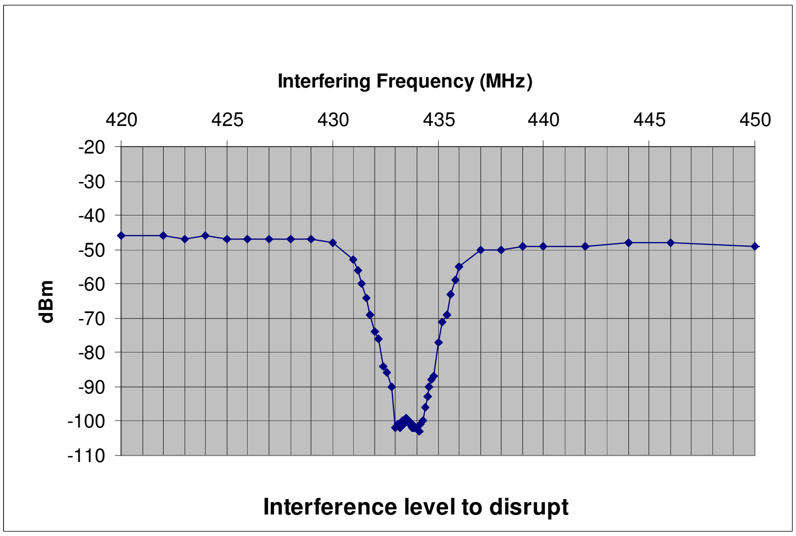

Figure 3

This receiver has no particular vices and a broad response, so generating the plot manually was not difficult. It is however, very revealing. For a 5 kbps receiver in an allocated band 1.7 MHz wide, this has a bandwidth of 2 to 4 MHz depending on the definition. We can also easily see that the limited out of band rejection is likely to cause trouble in some circumstances.

Example 2

The next tests were on a pair of more sophisticated Tx/Rx modules with RS232 data interfaces. This enabled the setting up of PC software to perform the pass/fail test. One transceiver was set up as a remote station and programmed to echo back a packet from the local station. The PC was thus able to perform a ping test by sending a simple message via its serial port and waiting for a reply. This method naturally gives us unwanted signal first testing. The test arrangement is shown in Figure 4.

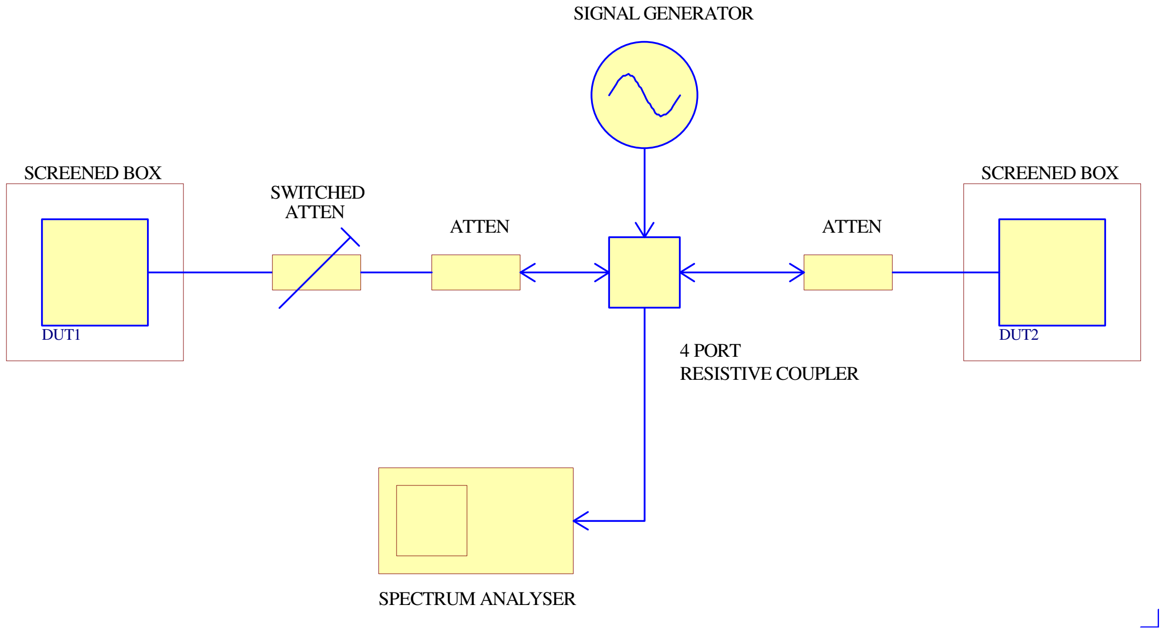

Figure 4

Once we have the pass/fail testing under computer control we can move on to automatic control of the unwanted signal. Various algorithms for the amplitude and frequency stepping of the signal generator were tried out, with the aim of getting a reliable but rapid finding of each point on the plot. Even so, it can be a lengthy process. The results in Figures 5 and 6 are assembled from a series of automated runs and some manual measurements.

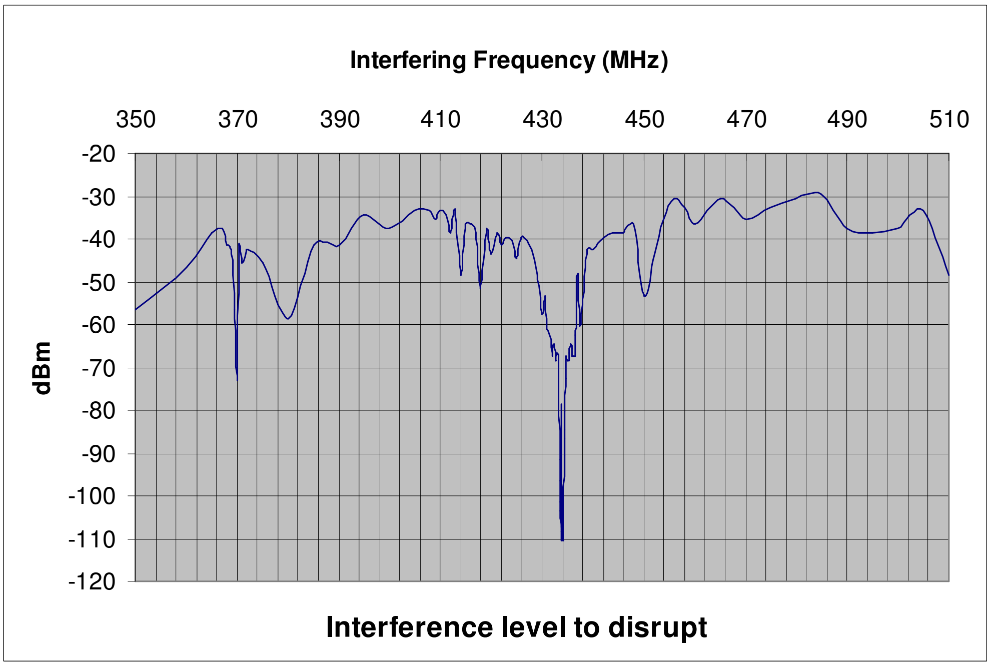

Figure 5

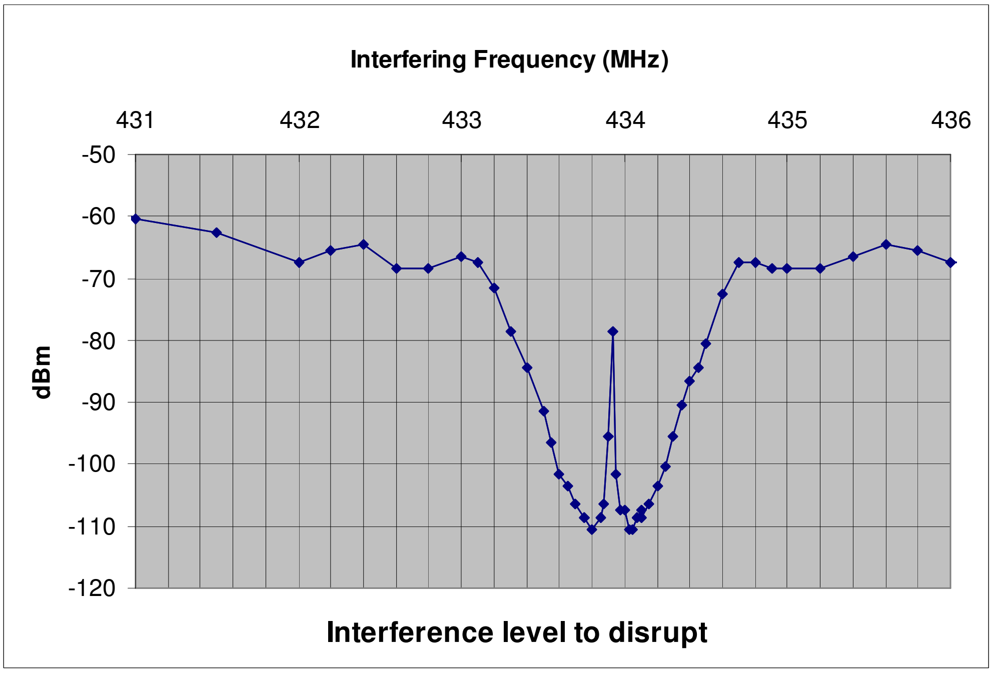

Figure 6

Again we have some interesting results, and ones that it would be difficult to describe using the traditional parameters. Out of band, the performance seems to be not so much a blocking level but a series of spurious responses. It would be necessary to make a more finely detailed survey before trying to interpret this further.

Closer in, we can see that the receiver is a low IF system, but with no image rejection. Furthermore, although this is a 20 kbps system sold as having 10 selectable frequencies, we can see that it is really only able to work with two or three distinct channels in the band.

What is noticeable though, is that at some frequencies, this device is even more susceptible to interference that the first example. It would appear that there was no significant RF filtering on this device, which was confirmed when the unit was opened up; it contained just a commercial RF chip directly connected to the antenna.

Example 3

What was apparent from the previous tests was the need for very fine frequency steps in the region of a spurious response. For the next tests the ideas of adaptive searches or manual intervention were dropped in favour of simply making fine steps all across the measurement, even if it did lead to runs of 6 hours or more.

The third example is a demonstration kit from Xemics Semiconductors. The test setup was similar to that of example 2, except that two units were set up to run their own ping-pong test. A pin on one of the demo boards gives a pulse for each successful round trip and this was monitored by a GP-IB connected scope. As before, a round trip test neatly satisfies the unwanted signal first requirement.

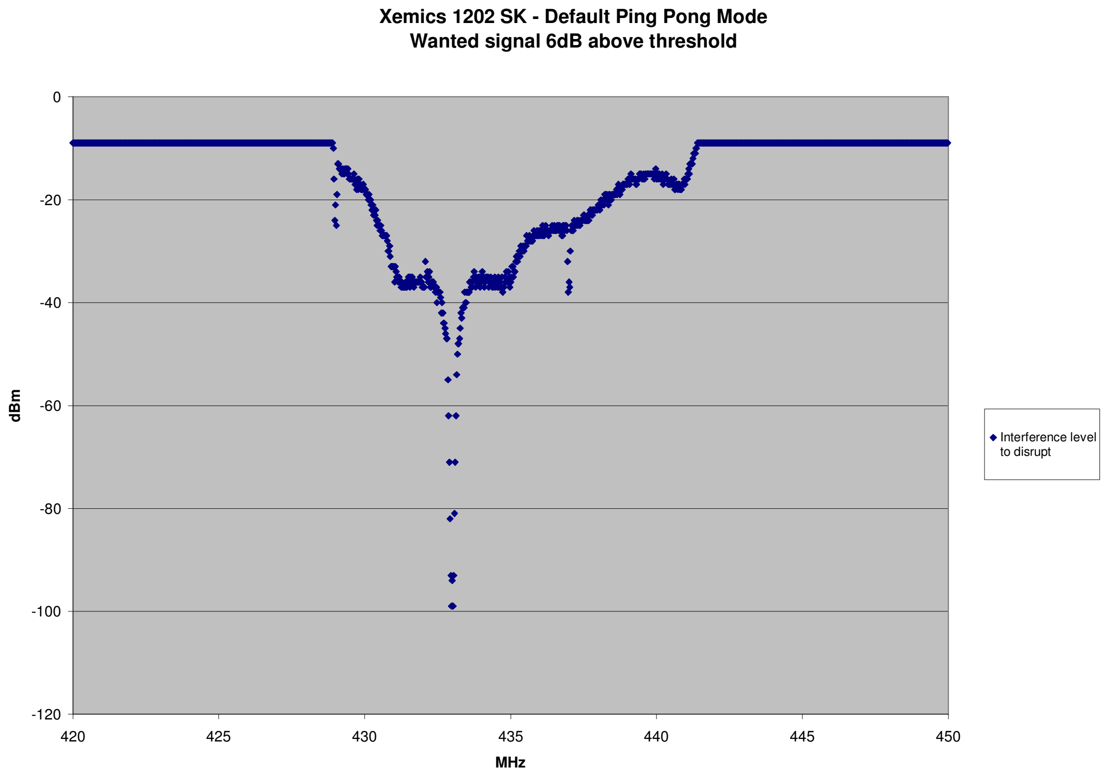

Figure 7

Figure 7 is the result of one automated run stepping every 25 kHz and Figure 8 is an assembly of several shorter runs.

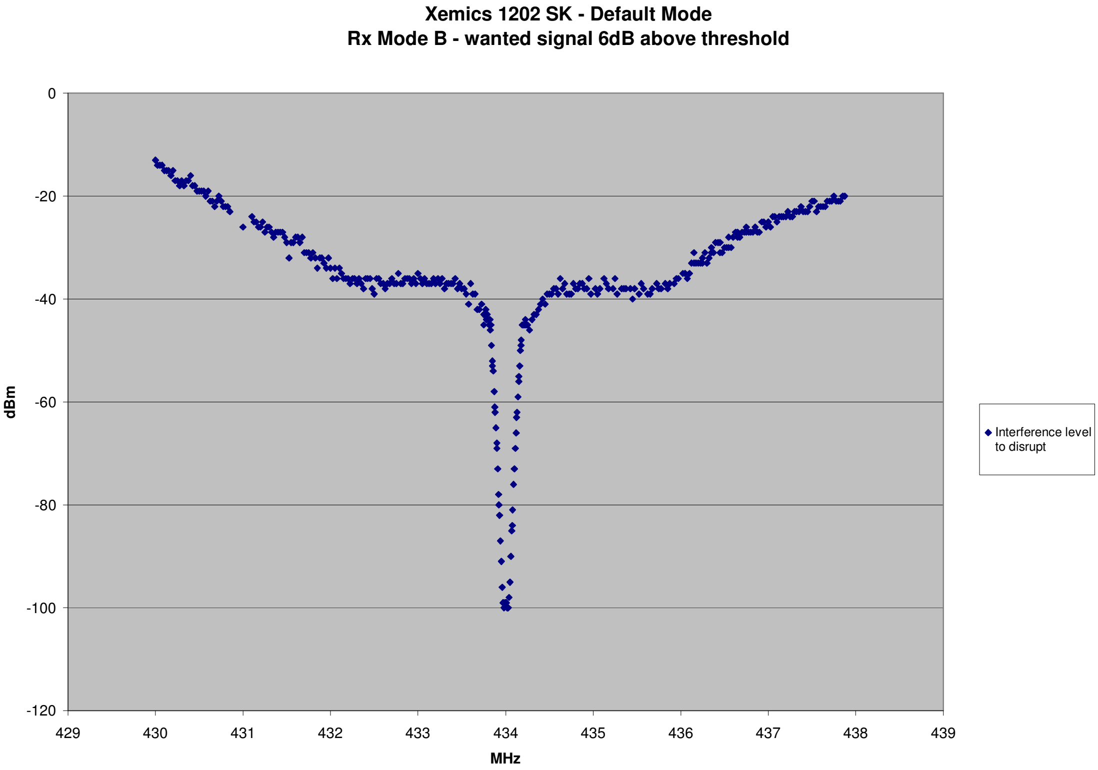

Figure 8

This time we are seeing the plots of a well behaved receiver. In Figure 8 as we move off either side of the co-channel response the rejection increases rapidly. It plateaus at a blocking level of –35 to –40 dB and this then improves further as the RF filtering takes effect.

The wide plot shows two low level spurious responses. These correspond with spurious products seen on the transmitter and so are probably synthesiser artefacts.

Conclusions

There are two stories here, the radio and the testing. For the radios, it can be seen that there is a huge variation in performance. One of the problems is that this market is characterised by unknowledgeable people buying and selling on price alone. Examples 1 and 2 are really only suitable for convenience devices such as replacements for IR remote controls. Indeed, it could be considered negligent to deploy them in any other application.

Example 3 is admittedly a later product and hopefully more typical of current practice. Considering that it is a transceiver with a component cost of about £4 it is actually quite good indeed. This could be used with confidence in the presence of in band and out of band signals.

As for the testing, constructing a well diagram is a very powerful way of analysing a receiver's performance. It presents a large amount of information in an easily understood way. For checking susceptibility to particular environments it is much more useful than a list of parameters such as blocking and spurious levels. It does, however, require a large number of measurements but these are easily automated. Particular points to consider are:

> The interface to the radio. Since we are not dealing with standard products or standard tests, this is likely to be a custom design. Each of the examples above has been different and has usually required writing a special module for the software.
>
> It may be tempting to use a bit error rate test instrument as the interface, but for these purposes a BER test is only valid for continuous streaming, such as broadcast links. For most other systems, representative testing requires short bursts and unwanted signal applied first.
>
> Frequency and amplitude steps. Eventually we settled on laboriously stepping through the frequency range in fine steps, smaller than the receiver final bandwidth, and restricting the adaptive searching to the amplitude axis. The software starts with an estimate and large step sizes, eventually homing in with 1 dB steps and "best of three" decision making.
>
> Defining test conditions. The exact nature of the unwanted signal often makes only a small difference, but it should be recorded each time. We use an 80% AM modulated carrier as that is a fairly destructive signal and corresponds with standard EMC testing. Similarly the wanted signal level is important and should be recorded. Suitable values are 6 to 20 dB above the minimum useable.
>
> Robust software. Test runs can be many hours and it is important to ensure that the software does not crash with internal overruns or trip up if the signal generator changes range or reaches its amplitude limit. We write the data to an Excel readable file as it is generated so that it is not lost if there is a crash.

Summary

A powerful means of testing and displaying the interference immunity performance of radio receivers has been shown. The results can be shown in a graphical form in a way that enables easy interpretation.

In the process it has been shown that the performance of low cost commercial products can range from extremely poor to surprisingly good.

Nick Long
Great Circle Design
Pine House, High Street, Wincanton, Somerset BA9 9JF UK

nick@gcd.co.uk Tel: 01963 33213 www.gcd.co.uk