TEKTRONIX 2465B A5 RANDOM CPU LOCKUP PROBLEM.

Dr. H.Holden. March. 2026.

I run a few 2465B scopes in my own workshop and in the past have written a number of articles about repairing these scopes and keeping them running.

About a week ago after one scope was powered for a few minutes an interesting event occurred. A number of the front panel LED's appeared to flicker for a moment and then went to a random pattern of selections. Once that had happened the panel was frozen with no response to any button presses. The scope display vanished, both the readout and the traces. The beam finder had no effect.

When I powered the scope off I noticed that there was a flash from the screen phosphor indicating that probably the CRT's high voltage support circuits were basically operational.

The scope was obviously still in its housing therefore I could not make any meaningful tests. I removed the scope from its housing and I powered the scope again. It booted apparently normally with no errors and didn't go into diagnostic mode and appeared to be working again. I left it running watching it like a Hawk, within an hour, the same thing happened, except this time it was a different pattern of panel LED's that were lit. One initial thing I did was physically inspect the A5 computer board, all looked normal.

At that point I performed the usual tests and checked the power supply. All voltages were normal. This did not surprise me, because it was only a little while back that I recapped the PSU and also wrote a detailed article on this topic:

https://www.worldphaco.com/uploads/2465BPSU.pdf

Unfortunately the fault appeared to go away for a whole day. Later it returned. While the scope was in the locked up condition I was able to make some tests with another scope around the CPU. I found that the small blue 10MHz clock module was running normally and the CPU was being clocked by half that rate on pin 39 and also the CPU was not in a Reset condition.

Of note, with the 2465B when it is running out of its housing, it is important to apply a Fan to U800. This is because the air flow across U800 is lower and that U800 is cooled primarily by convection. High level thermal cycling of U800 encourages delamination of the output transistors in its package. In my scopes I have fitted copper heat-sink that attached with a custom machined copper bush to U800's "TO-220 like" metal tab, this also helps, it is much better than just applying a heat-sink to U800's epoxy case.

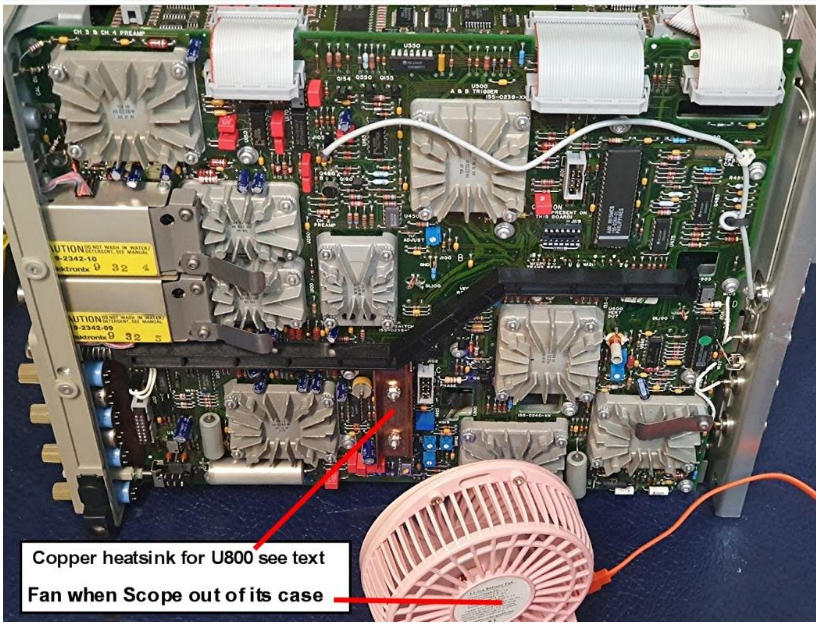

The CPU board (A5 Board) in my scope is the later version which uses the Dallas DS1225 non volatile battery backed up SRAM.

Fortunately this A5 board was a variant that had mostly Tantalum Capacitors rather than the physically leaky & destructive surface mount Electrolytic Capacitors of some of Tek's A5 boards. The two electrolytic capacitors it does have are a Radial type which stand off the board and not prone to leaking electrolyte.

The 2465B's A5 Board.

On testing, in the locked up condition initially, the CPU's R/W line was stuck in the read state and all of the Data lines were low and the address lines had a mixture of both high and low static states, nothing was changing state.

I found that line power cycling the scope again cleared the lockup condition and I had to wait some minutes to hours for the fault condition to return. I also wondered about resetting just the CPU, so next time it locked up, I transiently pulled the the CPU's reset pin low again with a resistor for current limiting. The lockup condition was cleared for a while at least.

Also every time the scope was power cycled (or the CPU alone reset by toggling its reset pin) it would recover from the locked condition and the scope, for a good while at least appeared to be working.

(However, I missed initially an idea for a test and only realised later, after I found the fault, there was something else I could have done which would have given the game away by testing the the scope in the working condition - see below)

Then, while studying the problem, in the locked up condition another version of the condition appeared:

This time despite the buttons and controls being locked out and having no CRT display the address lines all but one A9, were active and so were the data lines. In this case the CPU was executing code, but was stuck in a relentless loop that it could not escape from, but it did escape with a reset.

Also I noticed that the /IRQ line was continually active with a square wave at 310Hz indicating that some peripheral was probably trying acquire service but the CPU was unable to complete it. I decided that something had caused the CPU to jump into the weeds and either totally lock up or not be able to escape a program loop it had jumped into.

I started to formulate some ideas about what might be wrong.

In the past I have had a 2465B with a Tin Whisker issue. When Tin Whiskers form they can be intermittent and make an initial connection that fuses after contact with a terminal. When there is voltage between the two terminals that can source enough current to fuse the Whisker tip.

Though a Whisker was fairly high up on my suspect list, after about 6 or 7 episodes of CPU lockup, I found it more difficult to imagine that a Whisker (if present) transiently shorted out something like adjacent address or data lines and always disconnected immediately after that, so as to allow a successful reboot.

The **Catch 22** with intermittent CPU lockup scenarios:

Unless provision is made to detect and escape from it, if a transient event occurs that disrupts the program flow and the CPU vectors to some random memory sector, it can remain stuck with invalid code. In some cases, it can even get stuck in short loop executing some code if there is some valid code at that memory location to execute.

Jumping into the weeds is a fair analogy. For example corruption in the stack area of RAM can cause this sort of thing or a Jump table pointer area. The condition the CPU ends up in, is not the actual fault that caused it. Marginal ROM or RAM and possibly even the CPU or data buffer IC's could all be responsible as well as large power supply transients.

Though we all know the old saying "it is never the CPU at fault" because many technicians with computer faults tend to find if they knee-jerk react and blame the CPU and replace it, over 99% of the time the fault remains.

The abnormal electrical event of an intermittently failing component or IC, or a Tin whisker short, or data corruption anomaly may well have long passed by and not even be currently present after the lockup event. In the intermittent CPU lockup fault scenario, to go looking for an issue that has not occurred yet prior to a lockup has a very low yield and in this case to go looking for a cause after the lockup can also be equally difficult.

Therefore the question: Where is the problem that caused the lockup? It becomes a matter of statistical probability, rather than usual fault finding methods. The usual fault finding methods, for locating some intermittently faulty component, or pcb issue, involve both powered or un-powered diagnostic testing of the pcb and its components.

One other possibility may have been the clock randomly dropping out or having glitches. To exclude that, I deliberately manipulated the clock on and off as an experiment and found that in this circuit the CPU does not go into a lockup mode as a result. And on testing, in either the working mode or the lockup mode, the clock operation always appeared perfect. A clock issue started to fall down the suspect list, even though the small blue 10MHz Cmos/Xtal clock modules on these A5 boards have failed before.

Then there is the issue of possible RAM corruption. For example if the RAM becomes marginal or faulty and that part of the RAM is being used for the Stack or a return address or a jump table pointer, it can send the CPU to no-man's land somewhere else in memory.

I had initially only a little suspicion the RAM would be faulty because it was only programmed and verified within the last year and was perfectly normal then.

After making a number of measurements around the CPU, of the Data lines, Address lines, /IRQ line, Clock, Rest line, R/W line etc, in both the locked and the unlocked condition I decided to look at the RAM.

In the case of this particular scope, rather than running my FM16w08 FRAM upgrade, it was running a re-worked DS1225, with the Immortalisation method I devised see below.

This DS1225 module was known previously at least a working part. In my case I keep the RAM files and spare programmed RAMs for each of my 2465B scopes. I also have records of those files over 10 years or more and I can compare the bytes in the calibration area.

Previously, I had noted that the bytes corresponding to the scope's calibration constants are kept in the RAM over the address range in the RAM itself of 1E00h to 1F55h. And I had also noted that the byte pairs there remain invariable over time (after scope calibration) and use of the scope. Except for the first 11 Words (22 bytes) The reason for this is that these get altered by the DC balance procedure and those particular calibration constants in the calibration area change over time, but the rest of the calibration data does not and should never change, unless there is data corruption.

Therefore, as an experiment I took the Dallas DS1225 NVRAM out of the scope and fitted another known good programmed spare. The scope booted normally and I waited, to see if the intermittent lockup would occur.

In the meantime I decided to put the removed DS1225 in my GQ-4x reader/programmer to have a look at its file and compare that, especially in the address area of the calibration constants, with the file that it had been previously programmed with.

To my surprise all of the bytes in the calibration area which had a value where the second bit of their upper nibble was previously low, had all gone fixed high. Not only that, this problem affected the entire RAM's content, here is a snapshot of some of it:

What is wrong with this picture ?

```
00000000 3C 3E 60 20 60 20 60 20-60 20 70 20 70 20 70 20
00000010 70 20 3D F8 2D 2B 38 7E-68 2E 78 F4 6F FF 23 34
00000020 36 B3 23 28 35 B5 67 B1-25 73 25 73 7E 3F 21 20
00000030 20 25 60 20 A0 28 20 A0-A0 60 20 32 E0 20 61 25
00000040 A4 23 64 20 6A 2A 21 22-A8 34 21 20 20 2C A8 FF
00000050 A2 20 F8 A4 2B 2B 2B 23-25 3E 22 2A 78 60 2A 20
00000060 22 20 20 2B 2F 3F 20 20-20 20 20 2D 22 20 B0 A0
00000070 68 20 75 FA 20 20 B5 A9-30 E0 20 20 24 FD 20 20
00000080 6B 22 A2 28 65 23 2C 23-E0 22 A8 22 A2 28 20 E0
00000090 F3 20 AB 20 BE 3D 31 3F-AB 3F AE 3F A5 2D 66 22
000000A0 2D 3C F6 3F A7 3F AD 3C-B9 2C E2 22 E6 20 2D 20
000000B0 20 20 2C 31 34 63 3A 30-37 7C 22 A8 22 A2 2D F8
000000C0 22 2D 26 3B 2D 2B 22 E6-26 68 20 20 20 25 31 35
000000D0 67 20 3C 2F 20 20 28 30-E6 6F 29 20 22 AC 20 3E
000000E0 20 A0 26 B3 22 7F 20 3F-20 A0 25 B5 20 3B 20 20
000000F0 20 24 20 21 20 20 20 20-20 FF 23 34 23 28 20 20
00000100 20 20 20 20 E0 20 20 20-20 21 24 30 62 A4 22 3C
00000110 63 A8 B0 20 20 60 20 20-70 20 21 25 25 25 22 22
00000120 21 FF FF FF FF FC FC FF-21 20 25 23 20 25 FA 20
00000130 20 3F 20 20 20 25 B6 20-2A 20 20 20 E8 20 20 20
00000140 64 20 20 30 20 20 20 20-23 2C 20 20 20 20 20 20
00000150 23 E8 20 20 20 20 20 20-20 22 20 20 20 20 32 33
00000160 20 20 32 21 20 20 20 33-20 20 70 6F 2A 78 20 20
00000170 20 30 20 20 29 20 3B FD-FE 21 20 20 20 2D 2B 20
00000180 20 20 20 20 20 20 E8 20-21 29 20 21 27 AE 20 20
00000190 20 20 20 A3 20 20 22 7C-22 60 20 20 22 AF 2E 3B
000001A0 20 20 20 20 20 20 20 20-20 20 20 20 20 20 20 20
000001B0 20 20 20 20 21 BC 20 20-20 20 20 20 20 20 20 21
000001C0 20 36 20 27 23 EB 20 20-20 20 20 35 20 20 2A 23
000001D0 EC 20 2A F3 BA E2 20 20-20 20 20 20 20 20 20 20
000001E0 20 20 20 20 20 20 20 20-20 20 20 20 20 20 20 20
000001F0 3A B0 75 36 BF A1 64 F0-20 20 20 20 20 20 20 FD
00000200 20 21 63 20 22 FC 27 20-20 20 20 20 20 20 20 20
00000210 20 20 20 20 20 20 20 20-20 20 20 20 20 20 20 20
00000220 20 20 20 20 20 20 20 20-20 20 20 20 20 20 20 20
00000230 20 20 20 20 20 20 20 20-20 20 20 20 20 20 20 20
00000240 28 20 20 20 20 20 20 20-20 20 20 20 20 20 20 20
00000250 20 20 20 20 20 20 20 20-20 20 20 28 7B 23 3B
00000260 22 AA 2E 2C 2F 2F 2F 2F-2F 2F 2F 2F 2F 2F 2F 2F
```

All bytes that were once 00 had gone to 20 etc, every byte had a value of “2” added to its upper nibble. This pattern with that “bit stuck high” had occurred throughout the entire DS1225. Of course this did not affect the value of a byte where that bit was already high.

What surprised me somewhat was that the scope managed to boot to an “apparently normal” operating condition initially at all, with probably at least 50% of the bytes being abnormal. A correct write then read value in RAM at a location used by the CPU, would only get returned if the second bit in the upper nibble was already a high value for that byte.

After thinking about this issue for a while, it is likely that most of the code required to initially boot successfully is in ROM. It is only when temporary values held in the RAM that are responsible for maintaining integrity of the program flow such as the Stack gets corrupted or a return address, that the CPU then jumps into the weeds. There are plenty of other address locations where corrupt RAM data would have little effect, for example a change to another menu state or another measurement scale or some unused flag bit.

It was also then I realized I had missed a big diagnostic opportunity. Since the scope was “working” for long periods of sometimes hours before a lockup, if I had performed a calibration check on it, it would have been abnormal because of the grossly corrupted calibration constant section in the RAM and that would have alerted me to a RAM corruption issue right away. Things seem easy with hindsight.

Also, while I have seen a number of DS1225 chips with various defects, such as a few unreliable cells, this is the first case I have seen where the effect was to jam one bit in the upper nibble of the byte for all memory locations and cause an apparently intermittent lockup condition of the CPU.

I also discovered that the RAM socket had become intermittent, in that it was fine with the FRAM on the adapter board, but it was giving trouble with an actual perfectly normal DS1225. I discovered this at the end of the repair process. It is worth considering this issue, not just for the 2465B A5 computer board, but for any vintage computer, especially with RAM in sockets.

IC socket Dilemmas:

In the past the original Dallas DS1225 NVRAM had been removed from this A5 pcb and a socket fitted, in this case was a Dual Wipe type.

Considering all of the variables, why would it be better to choose a dual wipe socket over a machine pin type with the Gold inserts?

The answer depends on whether the computer in question, any computer, is still under repairs, be it the scope computer in this case or any vintage computer and if the other repairs are complete yet, or not.

The Dual Wipe socket is easily damaged. I think the modern ones are not made of the same springy Phosphor Bronze or Beryllium derivatives they once were and some might be plain Brass.

The adapter board previously used was fitted with FM16w08 FRAM. They were some of the smallest pin sized adapters I could find at the time. Thin pin adapters, which resemble IC pins, are available and those would have been a superior choice.

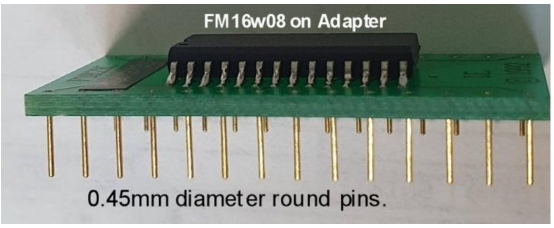

This FRAM & Adapter board was from a previous historical project where I fitted FRAM to the scope. After I did this and Web published the paper on it back in 2013, many 2465B scope owners followed suit and converted to FRAM:

https://www.worldphaco.com/uploads/TEKTRONIX_2465b_OSCILLOSCOPE_CALIBRATION___REPOWERING_THE_DS1225.pdf

You might imagine that the Dual Wipe socket would have put up with 0.45mm pins and later would have still been viable with the pins from a Dallas DS1225, but they were not.

Despite having started the FRAM solution to the DS1225/2465B scope dilemma, I had moved back to the Dallas DS1225 after devising a way to Immortalize them by replacing the battery with what amounted to a battery compartment. The battery could be changed at will. In the end I settled on a 10mm Milling tool as it tends to cut just a little over a 10 mm Diameter with the one procedure and provides satisfactory clearance for the battery:

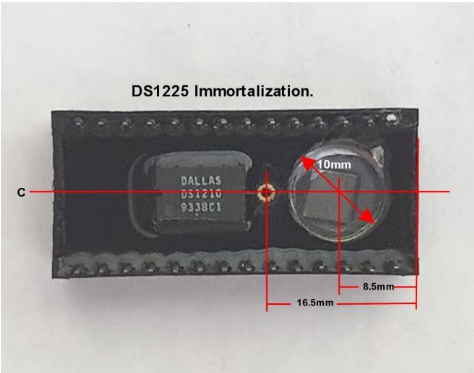

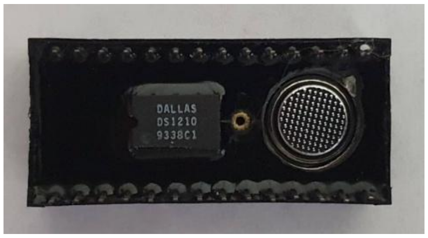

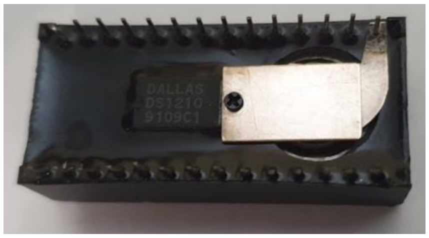

The details on how to execute this on the DS1225 are in my Web article here:

https://worldphaco.com/uploads/THE%20DALLAS%20DS1225-ReMastered..pdf

On re-inserting a normal programmed Dallas NVRAM (it has 0.27mm x 0.53mm pins) into the socket where the FRAM adapter had been, on one occasion one of the pins did not make contact. This indicated that the spring leaflets in the dual wipe socket do not have much mechanical "memory" at all and they were permanently deformed by nothing more than a 0.45mm round pin.

Many round pins on various adapters and headers are closer to 0.5mm to 0.55mm too. Of course in some kind of fault finding scenario, this sort of thing is a disaster, because faults start to compound as more work and investigations are done and the repair spirals down hill. This sort of thing can happen when RAM/ROM diagnostic adapters with 0.5mm round pins are inserted into vintage computers, such as the PET, to help with diagnosis and repairs. It can result in socket degradation. It is even worse there too because many of the sockets for the CPU/PIA/VIA & ROMs were only a single wipe type in the PET and they are less hardy than a dual wipe type.

It gets to the question: If you are fitting an IC socket or replacing one, which might be the better sort, Dual Wipe or round Machine pin? It is an interesting question**:**

DUAL WIPE:

Advantage; in the case of the dual wipe type, one highly favourable feature is that the pin geometry that passes into the pcb hole is super thin.

For example on the flat axis these pins are only around 0.15mm thick (and about 0.6 mm wide) This means that it makes removing the socket from a pcb a breeze, because there is plenty of air flow and molten solder for the sucker to work brilliantly, especially a single shot high volume sucker.

The solder is well cleared and not only that the pin does not get stuck to the upper pcb pad or the side wall of the hole. Once cleared of solder, the pin is easily rocked from side to side in the hole to make sure it is totally free on every pin, before the socket is lifted away and presto, no pcb damage.

Disadvantage; some types appear significantly damaged by an over sized pin of as little as 0.45mm thickness or diameter, making insertion of a real IC with a flat pin (typically around 0.25mm thick) unreliable later.

Machined Pin:

Advantage; There is a very favourable feature of the machine pin IC socket I call *"the two chance effect"* This is because of the anatomy of the four prong gold ferrule in the socket. Here is a photo of the typical 4 pronged insert from a machine pin socket:

In the case of fitting a standard IC pin to the machine pin socket (typical geometry 0.25mm x 0.5mm), only the two prongs on the 0.5mm wide axis of the IC pin are pushed apart to any extent and engage the IC pin side edges.

If a round pin is pushed into the socket in the range of 0.45mm (better) or even 0.5mm (at worst) the 4 prongs engage the pin surface equally. If that round pin is then removed, the socket’s insert still works for the flat IC pin because the two prongs on either side that are the main operational ones for the flat IC pin and are undamaged. The other two prongs played no significant role anyway for the 0.25mm to 0.3mm wide flat faces of the IC’s pin.

In other words, the machine pin socket is better able to tolerate both types of pins, and work for both and *switching between both*, especially if the round pin is not above 0.5mm in diameter. This is not the case for the dual wipe socket because they grab the IC pin over the flat axis which is only in the region of 0.25mm for a real IC pin.

Disadvantage; The machine pin socket is somewhat diabolical in one respect. It has a conical area where the diameter of the pin tapers down to the part of the pin that passes through the pcb’s holes.

What this does is that it often occludes the mouth of the plated through hole on the top surface of the pcb. That blocks the flow of air that allows the solder sucker to work effectively. And it really needs to work effectively, because the round pin occupies a greater cross sectional area in the hole, leaving more opportunity for the pin to remain stuck to the inner wall of the plated through hole, around some of its circumference.

These are the reasons why many people who try to remove machine pin sockets sometimes end up ripping pads/tracks & pcb plated through holes.

The above are some considerations of Dual wipe versus Machine Pin sockets, but people might think...why get worried about it, why would I need to replace a new IC socket that I just fitted?

Well the answer is that unless the board is fully repaired yet, you may end up damaging that socket with adapter boards for other IC's, or with some kind of diagnostic adapter later. In reality if the board is not fully fixed yet, the risk of socket damage remains through diagnostic interventions and use of various adapters.

One other disadvantage, though it is not too common, some people have reported that occasionally an IC's pin will become trapped in a machine pin socket and it won't release. This can happen for two reasons. One is if the IC pin is bent and rotated 45 degrees. Or there was a manufacturing error at the factory and the Gold insert was 45 degrees off axis (possibly some of the cheaper sockets might be more affected) When this happens the edges of the wider axis of the IC's pin are forced between two opposing slots in the 4 prong insert. The arrangement then becomes not dissimilar to a Chinese Finger Trap, and pulling on the pins causes the prongs to tighten down on the pin, preventing the pin's release.

Conclusion & Solution:

The above led me to the conclusion that overall the Machine Pin type of socket was intrinsically better and more resilient to a range of pin geometries, if only the problem of difficult de-soldering could be solved.

After a number of experiments with this problem, I found that the answer was simple. All that is required, when fitting the machine pin socket to the pcb, is to simply elevate it a little off the pcb 0.6mm to 0.7mm is plenty. This prevents occlusion of the mouth of the plated through holes and provides an annulus for the air & solder to flow. And it means that the solder sucker again becomes highly effective if the socket requires removal & replacement again, with a lot less risk to the pcb’s pads & plated through holes.

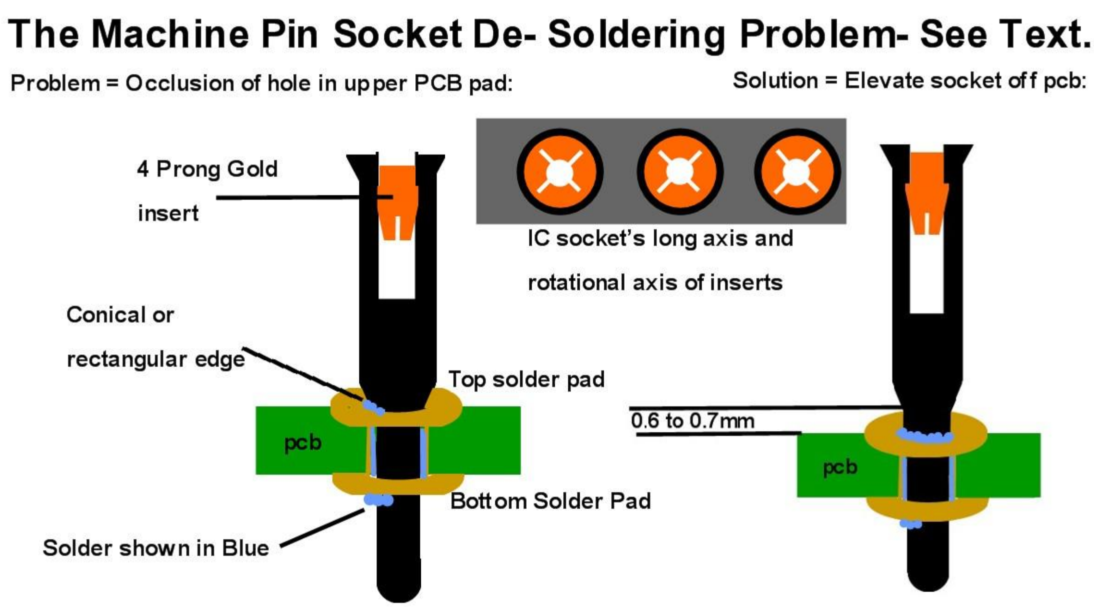

****