The ground plane: Lord of the Board

Tim Williams
Elmac Services
www.elmac.co.uk

Introduction

The way that a circuit is laid out on a PC board is critical to its EMC performance. At high frequency all tracks must be regarded as components in their own right, since the coupling of high frequency signals from one part of the circuit to another is mediated by their stray capacitance and inductance. Even if the equipment is a simple analogue circuit with no HF emissions sources, susceptibility to RF disturbances means that it must still be laid out with RF performance in mind. The schematic rarely if ever tells this side of the story; PC tracks appear as nothing more than lines on the diagram joining pins of lumped components such as ICs, resistors and capacitors.

But every track on a board exhibits its own partial self inductance, and mutual inductance and capacitance to other nearby areas of copper, and if these are uncontrolled then the performance of the board at high frequency – and that means the EMC performance of the circuit – is also uncontrolled. The simplest and most effective way of fixing the characteristics of all such tracks is to implement a ground plane on the board.

How a ground plane works

In this article, the term "ground plane" will be synonymous with "0V plane". That is, the reference (0V) for the power supply will be taken to be the same as "ground". This is not necessarily the same as "chassis ground" and there is no general rule (although there is a preference) as to whether the chassis should be connected to 0V or not. Since we are discussing just the PCB layout here, the issue will be largely avoided.

First of all, forget any notion of the ground plane being used for screening: it may do, but that is not its primary function. What you expect from it is to provide a high-frequency return path for every signal trace and power rail on the board.

Running a signal and its ground return close together – so that neighbouring currents flow in opposite directions – reduces the inductive impedance of the total path by a factor equivalent to the mutual inductance between the two halves:

Effective inductance of signal loop $=$ $2(L - M)$

where L is the self-inductance of each half of the loop, M is the mutual inductance between the halves
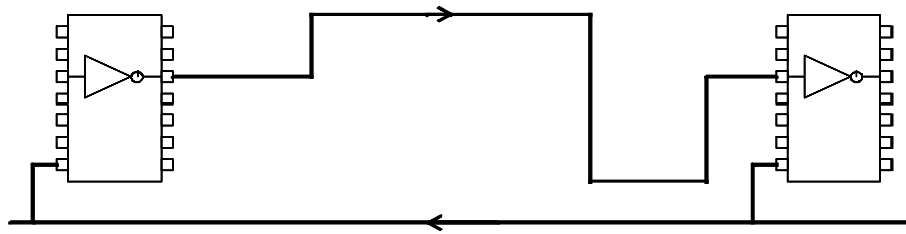
Bad practice: wide separation of signal and return, hence low M
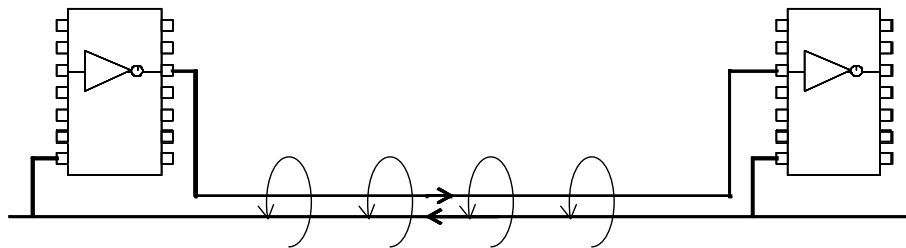
Good practice: close coupling of signal and return, M ≈ L

Figure 1   Separation of signal and return

This effect, which "kicks in" above a few tens to hundreds of kHz, is critical to understanding PCB track routing at radio frequencies: you need to control the 0V current return path. Keeping the signal and 0V paths adjacent along their length – and therefore maximizing their mutual inductance – will ensure not only minimum coupling with the magnetic fields around the PCB, but also minimum impedance of the return path and therefore minimum noise voltage developed along it. This will keep down the "noisiness" of the whole board, which is mostly what is responsible for common mode RF emissions. It will also improve its susceptibility to external interference, which is created by the conversion of incoming common mode disturbances to differential mode within the circuit.

Now, it's entirely possible to do this by assigning a 0V return track to each signal and running the two next to each other, or on adjacent layers, throughout the board. This is of course somewhat tiresome, not to say awkward, in a board with hundreds of such paths, and it isn't necessary. If you provide a 0V (ground) plane as a continuous layer adjacent to each track layer, then the job is done for you. The return current for each high frequency signal track automatically finds its own preferential path through the plane, which will be that path which is directly underneath the signal: the geometry ensures that this path has the least enclosed loop area and therefore the least total path inductance. All you have to do is make sure that the component 0V connections, which carry the return path through, have a short and direct connection to the 0V plane, which is easy; and that the 0V plane remains unbroken along the length of any of these paths, which despite also being easy on a multi-layer board still occasionally suffers from poor implementation.
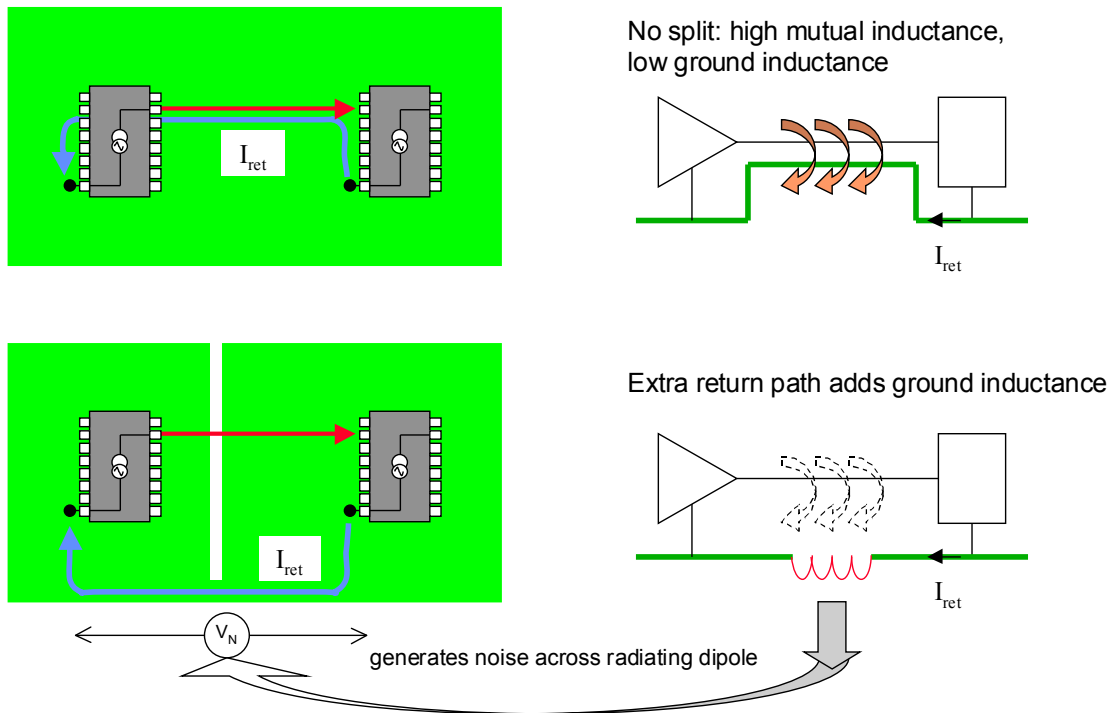
Ground planes in the xy direction: breaks, moats and joins

The effect of a slot across the current flow is to divert the return current from its optimum path and increase the 0V return path inductance. There are two EMC problems caused by extra inductance here:

- one is that any high frequency current passing through this inductance will develop a voltage across it. This voltage then appears across the two parts of the plane, each of which has a high capacitance to the environment due to its surface area. On its own this structure forms an efficient radiating dipole, which is then enhanced if cables are attached to the board;

- the other is that the inductance itself represents increased magnetic coupling to the environment through the enclosed loop area.

Therefore either there should be no split within the plane to add this extra inductance, or high frequency currents should be prevented from flowing between the two separated segments. This can be achieved through careful layout design for emissions control, by not running any tracks carrying such current across the gap. But it doesn't address the problem of immunity to incoming RF fields. The magnetic and electric coupling is still present for these fields, and so they will develop voltages between the two ground segments which can create circuit susceptibilities.

Figure 2   The effect of a discontinuity on return current

No split: high mutual inductance, low ground inductance

Extra return path adds ground inductance

generates noise across radiating dipole

There is a convention, encouraged by many device application notes, to create different 0V planes for different parts of the circuit: for instance, analogue 0V and digital 0V, which are then linked at one point on the board, typically at an A-D converter. The purpose of this trick is to prevent digital noise currents from flowing in the analogue circuit and corrupting it, that is, it relates to internal EMC. Unfortunately it creates problems for external EMC, and introduces difficult questions that could be avoided by a single system-wide 0V plane:

- where should you join the planes when there are multiple analogue-digital interfaces?

- what about power returns that are common to both analogue and digital?

- where do you run signals that must cross the break between the analogue and digital sections? Any of these signals will be exposed to incoming interference that will appear between the planes, or will develop such interference themselves.

- if you are going to bond the 0V planes at RF to the chassis of the product, which is always a recommended practice to minimise common mode coupling into or out of the PCB, how do you choose which of the planes should be bonded?
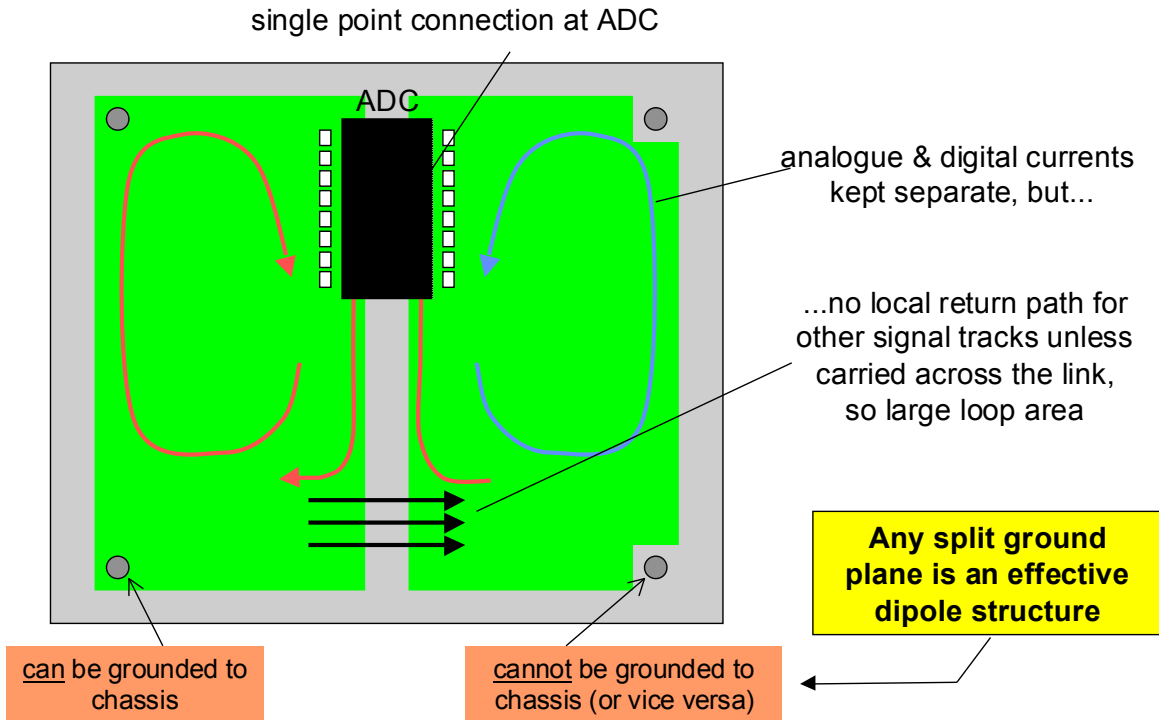
single point connection at ADC

ADC

analogue & digital currents kept separate, but...

...no local return path for other signal tracks unless carried across the link, so large loop area

Any split ground plane is an effective dipole structure

can be grounded to chassis

cannot be grounded to chassis (or vice versa)

Figure 3 The problems of split ground planes

For all these reasons, a preferable way to proceed is to use a single, system-wide 0V plane and carefully lay out the board so that digital noise is effectively segregated from low-level analogue circuits. Also, lay out the analogue circuits so that the critical nodes are subject as little as possible to ground noise voltages, that is, keep sensitive signals free of common ground impedance. Avoid as far as possible any discontinuities in this plane and make sure that you do not run any critical (noisy or low-level) tracks across any discontinuities that remain, such as slots formed by close-packed vias, or near to the edges of the plane. There is no merit in putting "moats" in this 0V plane.

In the centre of a plane inductive effects are minimal and the plane impedance is dominated by resistance, but towards the edges fringing magnetic fields are created and the inductive impedance of the plane rises. A good rule of thumb is to keep critical tracks further away than 10·h from the edge of the plane or the edge of any discontinuity, where h is the separation distance in the z-direction between the layer that carries the track and the layer that carries the plane.

This can also apply to power plane layers placed against a 0V plane. If the power and 0V planes were to have exactly co-located edges, the transmission line formed by the two planes would be terminated in a sharp open circuit and the problem of transmission line resonances would be potentially severe, leading to a high degree of coupling to and from

the PCB at the resonant frequencies. Staggering the edges of the planes softens the impedance mismatch and reduces the Q of the resonance. In fact, in contrast to the recommendation not to split the 0V plane, it is a good idea to split power planes into several smaller planes each of which supplies a particular segment of the circuit, and to decouple these from each other with small ferrite chip inductors.

The principle exception to the "no slots or moats" rule is where there is a good reason to *avoid* particular circuits coupling to the main 0V. Almost always, this is when such circuits are galvanically isolated: interface circuits isolated using opto-couplers, for instance, or data circuits using transformer isolation such as Ethernet. In these cases, there should be a window in the main ground plane within which all such isolated circuitry should sit. The edge of the 0V/ground plane should come up to halfway across the isolation barrier. Such isolated circuits need to be laid out as a unit so that there is no compromise of the isolation barrier, or the edge of the ground plane, by other circuits.
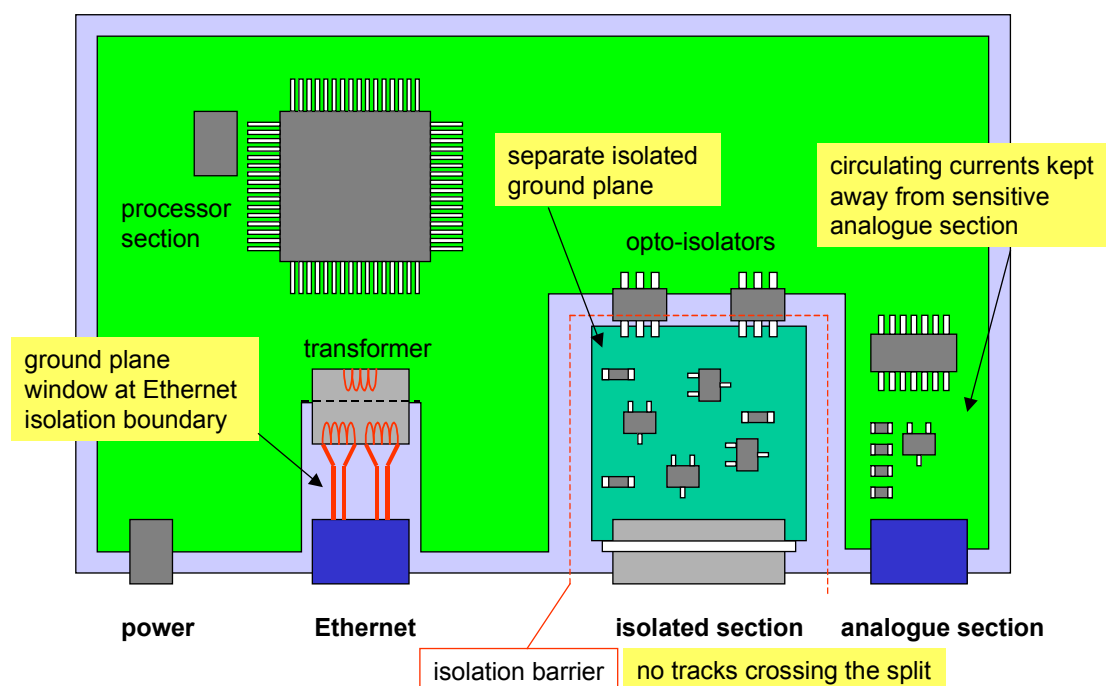

Figure 4   Uses of ground plane windows

Simpler circuits may use double-sided boards on which one side is devoted to a ground plane and the other side to all tracking, and the components. A good, solid plane can be implemented with some effort, but almost invariably some tracking has to be carried on the plane side, resulting in some apertures in the plane. The rule here is emphatically not to split the plane more than is absolutely necessary: keep any such tracks as short as possible and only use them as links, not to travel long distances. The resulting geometry of the cut plane should be analysed to check the effect of the cuts on return current flows under critical tracks.

Ground planes in the z direction: layer stack-up

Successful implementation of a 0V plane on a PCB requires careful attention to the layer stack-up, that is, the order in which the different tracking purposes are assigned to different layers. In multilayer (4 layers or more) configurations, every signal layer should be adjacent to a 0V or power plane layer. Also, power and 0V planes should be on

adjacent layers, possibly with deliberately thin separation, to take advantage of the interlayer capacitance for high frequency decoupling. Critical tracks, for instance those carrying high di/dt signals such as clocks, should be routed adjacent to 0V rather than power planes. Such tracks should also not jump through vias from one 0V plane layer to another, unless the 0V layers are tied together with vias near that point.

For boards with more than 6 layers, multiple 0V plane layers are needed so that tracks and power plane layers can always be located next to a 0V plane, but it is also essential to tie these together with grounding vias at small intervals – intended connections to 0V will provide this automatically, but in areas of the board where this doesn't happen, a 1cm spacing of such "stitching" vias will assure good performance of the planes up to about 2GHz.

Some consideration needs to be given to the question of whether 0V planes on different layers should or shouldn't overlap. Generally, a common geometry on all layers is easiest to work with and avoids the question. If there is a need for some reason to have partial overlaps, the principal consequence is that the inter-plane capacitance may create unexpected high frequency resonances, and the more complicated the structure, the harder it becomes to anticipate and control such effects. But don't assume they are negligible: as a simple example, a 2cm^2 overlap with a separation distance of 0.3mm in FR4 fibreglass laminate will have a capacitance of about 26pF; with 2nH of inductance in parallel (possible, with lack of attention to other details) this will resonate at 698MHz, quite capable of causing trouble in EMC tests.

One consequence of the dependence on layer separation is that the detail of the vertical sandwich construction of the PCB can be as important for its EMC as is the layout pattern. The separation is determined by the thickness of the pre-preg and core materials used in the build of the bare board. The danger is that this is left to the PCB manufacturer and is not specified by the circuit designer, or anyone else. As a result you end up with the default thicknesses used by a particular board supplier, which might be perfectly adequate for the EMC performance of the product and so is never questioned; but if another board supplier is chosen during the product life cycle, it is entirely possible that a different set of thicknesses could be used which result in changed performance. To guard against this, make sure that you specify layer thicknesses in the PCB drawing, if necessary checking with your preferred supplier what their defaults will be.

Constant impedance tracks over the ground plane

It is commonplace to find that some tracks, because they carry very high frequencies or are routed long distances, have to be treated as transmission lines so that reflections along their length are minimised. This contributes both to good external EMC in the form of lower emissions, and to good signal integrity, improving the reliability of the circuit. This means that the tracks have to avoid discontinuities in their linear and cross-sectional geometry. This in turn requires that

- they should only run in one layer, with no vias except at their ends

- this layer should be adjacent to one (microstrip) or two (stripline) continuous 0V planes

- the track cross-section dimensions and dielectric constant of the layer should be defined and controlled throughout the layer to give the correct constant transmission

line impedance: in practice, impedances between 50 and 100 ohms are easiest to achieve

- changes in direction of these tracks should be gradual rather than abrupt

This then requires that at least one layer on a PCB is devoted to these "constant impedance" tracks, and at least one adjacent ground plane layer (completely unbroken) is essential to maintain their characteristics.

Summary

The ground plane has enough functions on any electronic circuit board to ensure that at least one layer needs to be devoted to it alone. For all multilayer boards, it can be incorporated as a matter of course. Even for low cost double sided boards, with care a ground plane can be implemented on one side, provided that the rules about avoiding or dealing with breaks are understood. When such a plane is implemented properly, the external EMC of the board is automatically improved. Even simple analogue circuits are susceptible to exposure from radiated RF fields in the environment, and any designer who consciously avoids putting in a ground plane, on the grounds of cost, will find that satisfactory EMC performance will need greater cost and effort than if the plane had been implemented from the start.

*Tim Williams is a consultant with Elmac Services and is the author of **EMC for Product Designers** (4th Edition, Elsevier 2006)*

Copies of this paper are freely downloadable from www.elmac.co.uk

Originally published in EMC Journal Issue 72 September 2007

Elmac Services, Wareham, Dorset: phone 01929 558279