Splitting Planes

For Speed and Power

Douglas Brooks

Designers are sometimes confused about the question of how to deal with multiple power planes on printed circuit boards, and how to deal with their separation. We often talk about the horizontal separation between planes as a “split” and there are some design rules that may apply when designing in the area of splits.

But before we talk about how to deal with splits between planes, it is instructive to talk about why we have separate power planes in the first place – to accommodate separate power supplies. And one reason we might have separate power supplies is because we might have different voltages in our circuits. So the issue of separate voltages might be a reasonable place to start.

Power Supply Voltages:

On any given board, or in any given application, we may find DC power supply voltages of 5V, 3.3V, 2.5V, 1.8V, and even 1.2V. Older designs have included 12V supplies, and certain types of circuits may have special design requirements up to thousands of volts.

There are a variety reasons why these different voltages exist. One reason relates simply to the semiconductor technology used in the fabrication of the IC making up the circuit. Newer technologies and newer physical semiconductor structures often have inherently lower saturation and switching voltages, leading to lower overall supply requirements for their operation.

A primary driving force for lower voltage relates to the power required for the circuit’s operation. If the same circuit function can be achieved with a lower voltage supply, a power saving is usually achieved. Power is the product of voltage times current. Therefore, the power reduction can be directly proportional to the voltage reduction. But if the voltage reduction ***also*** leads to a lower ***current*** requirement (e.g. through Ohm’s law), the power reduction can be proportional to the ***square*** of the voltage reduction. This has obvious implications for battery operated circuits and for the thermal management of larger, more complex circuits.

Another driving force for lower operating voltages is rise time. If, for example, a technology can switch (dV/dt) a 5 V. signal in 2 nsecs, perhaps it can be modified to switch a 2.5 V signal in half that time, with no other technological changes required.

Finally, different circuit functions may require special voltages. These functions might include such things as transmitter stages, high voltages required by CRTs and laser printers, switching requirements for electromagnetic relays and machine tools, and audio speaker requirements.

Different voltages obviously require separate regulation circuits. This factor would define a ***minimum*** number of regulated power supplies, one for each individual voltage. But there may be additional, multiple regulated supply sections for any given voltage. For example, we may have different regulated supplies for the analog and digital sections of a board, even if they both have the same overall power supply voltage requirement.

We may also have different regulated supplies for different or sequential ***stages*** of the same circuit. For example, if we have different circuit paths for the individual R, G, and B components of a video signal, each path might have its own regulated power supply, even though the individual power requirements are identical. In extreme cases, some engineers design regulated supplies for ***each stage*** of a signal path through a circuit.

The reason for multiple regulated supplies is usually noise isolation. All signals flow in a closed loop. For every signal flowing down a trace, there is a return signal coming back. The return signal is usually on the plane closest to the trace, and (at least in fast rise time circuits) is usually as close as possible to the trace.

Howard Johnson illustrates this very nicely in his book. **Figure 1**, taken from Howard’s book, shows the current distribution of a return signal on the reference plane under a trace.

Now here is a point that is important to understand. Current is the flow of electrons. Figure 1 shows the variation in current density near a trace. It necessarily, therefore, illustrates the “electron density” in the same region. Electrons have charge. Charge density is voltage. Therefore, there is ***also*** a “voltage density” (actually a voltage *gradient*) that occurs near a trace.

Variations in voltage gradients can occur for several reasons. One is high frequency return current, as suggested in Figure 1. Another is simply the flow across the plane of current from the regulated power supply required for charging (and discharging) the various bypass (and planar) capacitances associated with the circuit. Voltage gradients across

This article appeared in Printed Circuit Design, a CMP Media publication, December, 2000

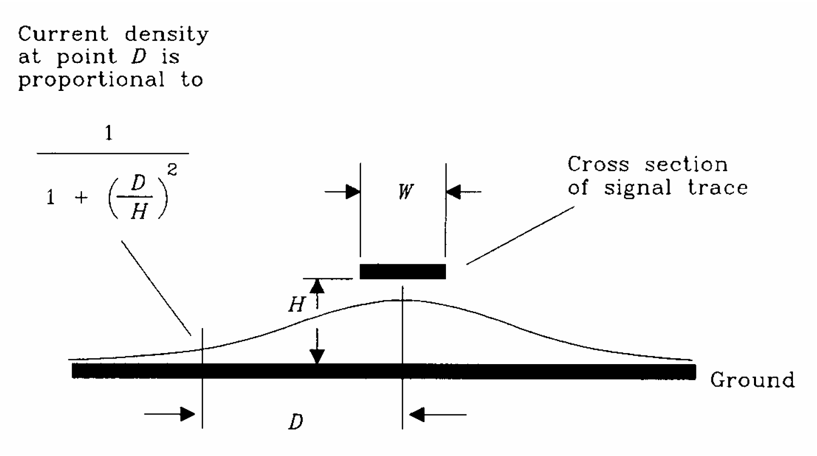

Figure 1
Current distribution of a return signal on the plane under the signal trace.

planes of as much as 250 mV are not at all uncommon on circuit boards.

These (changing) voltage gradients constitute ***noise***. And it is instructive to note that this noise is ***not*** coming from somewhere totally outside our circuit. In fact, it is generated by the circuit itself! One guideline for good power supply management is not to focus on preventing noise on the planes from getting into our circuits. It is, instead, to keep the noise that is generated by our circuits from getting onto the planes in the first place!

One reason engineers use different regulated power supply regions for different circuits and different stages of the same circuit is to try to prevent noise generated by one (part of a) circuit to interact with, and interfere with, signals in another (part of the) circuit.

Need for Power Planes:

Several articles have been published in PC Design Magazine in the past regarding the importance of power and ground planes in high-speed design applications. Here is a summary of some of the most important reasons why planes are important.

Impedance control: If we want to control trace impedance as a strategy for the control of reflections (using proper trace termination techniques), then good, solid, continuous planes are almost always required. It is ***very*** difficult to control trace impedance without the use of planes.

Loop Areas: Loop area can be visualized as the area defined by the path of the signal (traveling down a trace) and its return current. When the return signal is on the plane immediately under the trace, loop area is minimized. Since EMI is directly related to loop area, EMI is minimized when good, solid, continuous planes exist under traces.

Crosstalk: The two most practical ways to control crosstalk are (a) separation between traces and (b) closeness of the traces to their reference planes. Crosstalk is inversely proportional to the square of the distance between the traces and their reference planes.

Planar Capacitance: The capacitance formed by the proximity of two planes placed close together can be very important and beneficial in circuit decoupling at very high frequencies, where bypass capacitors and their associated mounting and lead inductance begin to have problems. And planar capacitance can be effective in controlling EMI radiations caused by both differential mode and common mode noise signals.

Strategies for Power Planes:

For all these reasons, the use of planes is very important and beneficial in PCB design. But then some relevant questions are, “How many planes should I use, what should be on them, and where should they be placed?”

For example, every different power supply voltage is typically distributed on its own plane. It is logical that each different regulated supply of the same voltage also gets its own plane; otherwise some of the regulated supply sections would simply be shorted out! But very often all the different regulated supplies are referenced to the same voltage potential – zero volts. Is it possible, then, to have only one ground plane (at zero volts) that can service all the individual regulated power supplies? Or do we need separate a ground plane for each individual regulated supply?

In looking at the question of whether each power plane needs its own, separate individual zero-voltage reference plane, we have to look at ***why*** the separate power supply exists at all. For example, suppose the overriding issue is power dissipation. Then, it may be perfectly fine for two supplies to reference the same ground plane. In some cases, a circuit simply involves two types of ICs with different supply requirements, and a single reference (zero-voltage) plane may work well here, too. And, if we have mixed signals (as in an A/D circuit) but all the circuits are known to be quiet, so noise is not a concern, a single reference plane may be adequate.

But, as noted above, one reason for using different regulated power supplies is for noise control – for example keeping transmitter noise out of a receiver section or keeping digital noise out of an analog section. And, ***all*** the reasons for using planes have a single common denominator – noise control. So it is almost axiomatic in high-speed designs that noise control is the predominant design issue, and our design strategies are undertaken with noise control as an important focus.

Now, if circuit noise is our primary focus, then we almost always need to define separate reference or ground (zero-voltage) planes for each individual regulated supply section. Consider the implications for not doing so. Assume, in a transmitter section, we send a signal down a trace, as shown in Figure 1. The return signal, while primarily underneath the trace, extends for some distance beyond the edges of the trace. If any receiver circuitry is anywhere near this return gradient, some of this transmitter noise may couple into the receiver section. In addition, there are current gradients that flow from the transmitter regulated power supply across the ground plane while charging (and discharging) any bypass capacitors that may exist. These currents may also couple into receiver circuitry. The whole point in having a separate receiver power supply section is to try to ***isolate*** the two sections. A single ground (reference plane) tends to work against this objective. So, good design practice is to have separate ground planes for the receiver and transmitter sections.

Similarly, a digital signal flowing down a trace generates a return gradient on the reference plane under the trace. Such a signal may cause a noise signal to be injected into a nearby analog trace. Again, for this reason, good design practice dictates using separate analog and digital ground planes.

In fact, anytime we use separate regulated supplies for noise control purposes, it is good design strategy to use a separate reference plane for each one, even though all the reference planes are at the same nominal voltage potential – zero Volts. Not doing so tends to defeat the very purpose for having separate regulated supplies in the first place.

Summary: Separate regulated power supply stages can exist for a variety of reasons. Very often they exist for the purpose of controlling circuit noise. Control of circuit noise almost always requires the use of planes. And if each different regulated power supply section requires its own separate distribution plane, then each should each its own individual (zero-voltage) reference plane, also.

Some Design Rules:

Now, based on the above, let's assume we have several regulated power supplies (associated with separate circuits on the board), each distributed on its own plane, and each with its own, identifiable, separate reference (zero-voltage) plane. What PCB design strategies and rules are appropriate?

Connecting Reference Planes Together: Ultimately, almost all regulated power supplies reference to the same thing, commonly a zero reference (zero-voltage) point in the system. Typically, if one puts an Ohmmeter between the various reference planes, we find that they are all connected together. It is usually important, however, that when reference planes are connected together, the connection is made at a ***single*** point.

Suppose, for example, this were not true. Suppose we had both a digital and an analog section on our board, and the digital and analog reference planes were connected at two (or more) points. There are conditions under which certain return signals ***could*** travel on ***both*** planes, (not *just* on the plane under the trace) thereby negating the vary separation we were trying to accomplish in the first place. Even worse, it is possible (under certain conditions) for noise signals to circulate in a loop around through both reference planes, crossing between them at the two connection points. Such current loops are called "ground loops." Their origin (cause) is often obscure, but their effects usually are not! The effects include noise problems, EMI radiation problems, and in extreme cases power dissipation and heat problems.

Control over ground loops is relatively simple. If there is only a single connection point for the reference planes, there is no loop over which a signal can travel.

But an interesting question is, "Where should that single point be?" On some systems, particularly where there are mother and daughter boards, the planes are separately routed back in a "star" fashion to a single point, usually at the primary power supply for the system. If there are multiple regulated power "islands" on a board, these might be connected to each other at single points with zero-Ohm resistors (really just a jumper) or with ferrite beads. The beads, of course, are used to block higher frequency components while still allowing a DC connection between the planes. Some argue, however, that if we have been really effective in isolating our power requirements, there are no higher frequency components to filter, and therefore a zero-Ohm resistor is all that is necessary.

In A/D circuits, it is common to have digital power and ground planes on the digital side of the IC and analog power and ground planes on the analog side of the IC. We often connect the analog and digital grounds together at a single point directly under the IC (or at least very close to the IC) with zero Ohm resistors or with ferrite beads.

Overlapping Planes: If we have separate regulated power supplies with their own reference planes, it is good design practice not to let unrelated portions of the planes overlap. For example, don't let a portion of an analog power plane overlap a portion of a digital ground plane (see **Figure 2**). Remember, a capacitor in its simplest form is simply two conducting surfaces separated by a dielectric. The area over which two planes overlap forms a small capacitor. It may be a ***very*** small capacitor, to be sure. Nevertheless, ***any*** capacitance provides a path over which noise may travel from one regulated supply to another, working to defeat the very purpose for which the separation existed in the first place!

An important part of the PCB placement process is to place components in such a way that their common regulated supplies (and grounds) can be efficiently grouped together.

Decoupling to Wrong Plane! We use bypass capacitors to decouple our circuits – i.e. to connect the power and ground planes together (from an AC standpoint) at specific points on the board. It is probably obvious that we don't want to drop a bypass capacitor from one power plane to an unrelated reference (ground) plane. Again, the reason is that we can (and almost certainly will, in this case) couple noise from one regulated supply section into the other. Unfortunately, this mistake can occur accidentally, sometimes fairly easily. As designers, we must be careful to check our engineers' net lists to ensure that this mistake hasn't happened. Even worse, we must guard against making this mistake ourselves. It is ***very*** embarrassing when this happens! There is simply no good answer to the question, "Why did you connect it ***that*** way?"

Signals Crossing Separations: Remember that signals reference to their (usually) nearest plane, be it power or ground. If we have a separation between two reference planes, we never want to route a trace across that separation. **Figure 3** illustrates an example of routing a digital trace across an analog plane. There are three primary problems that can result from doing this, all of them bad! They parallel three of the reasons for using planes in the first place (see above.)

(1) Good impedance control requires continuous control over geometry, and a continuous return path underneath the trace. If the trace crosses over the boundary between two planes, the return signal cannot "jump" the gap. (Don't assume you can solve this problem with a bypass cap. Doing so would couple two unrelated planes together!) This will cause an impedance discontinuity and a reflection, and therefore a potential noise problem at that point.

(2) If the return signal can't "jump" the gap, it must find some other path. This almost certainly increases the loop area for the signal and therefore the potential for EMI.

(3) Suppose two traces cross a separation between two planes. Since their return currents cannot "jump" the gap, these return currents must find another path. Even though the signal traces ***are*** separated from each other, their return signal paths might ***not*** be separated, and their return signals might "crosstalk". Thus, when signals cross plane splits, crosstalk may result even though there is no apparent cause or reason. This type of crosstalk can be ***very*** hard to diagnose!

The primary way of avoiding the problem of traces crossing splits in planes is to be careful in the layout and placement stages of the design. Group circuits by regulated supply voltage and then by signal flow. If a particular trace must extend into another regulated supply area, perhaps little "islands" or "peninsulas" (on ***both*** the power and ground planes) can be created for the traces involved.

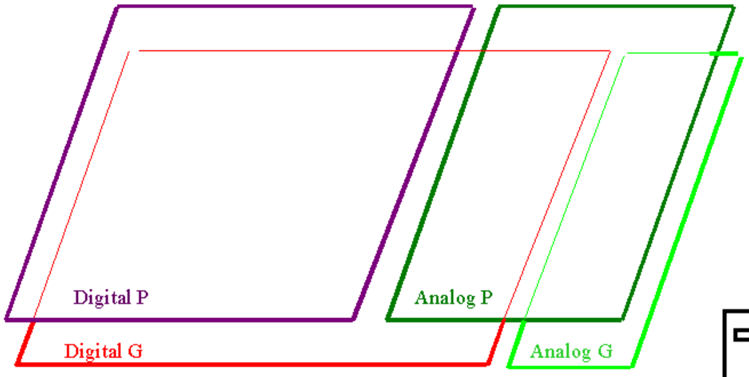

Figure 2
Signal coupling may occur if non-related planes are allowed to overlap. Note, the capacitor symbols represent capacitive coupling between the planes

Figure 3
Do not allow signal traces to cross plane splits or to travel over unrelated planes.

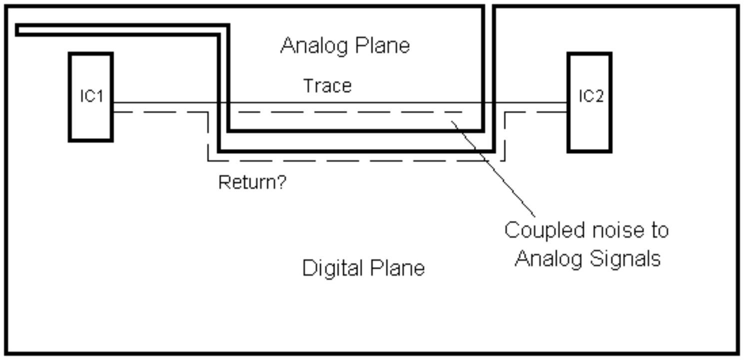

If there is no suitable way that can be found for routing traces without violating this design principle, then the circuit design engineer should be consulted. The engineer is the one responsible for deciding whether the design guideline can be relaxed in this particular case, whether the layout should be changed, or whether (in extreme cases) a different component selection is called for.

Conclusion:

It is common to have different regulated supply voltages on a circuit board. If circuit noise control is not a major issue, then it is possible for all these supplies to reference a common (zero-volt) "ground". If circuit noise – especially high-speed circuit noise – is an issue, however, then the regulated supply voltages are commonly distributed on their own individual planes. In this case it is normal for each regulated supply "plane" to have it own individual (zero-voltage) reference plane. These planes are usually constructed as "plane pairs" on the board. The boundaries defining individual planes are often called "splits."

Certain, relatively simple design rules are followed when we have split planes:

1. Don't allow non-related plane areas to overlap.
2. Connect reference planes together at only a single point.
3. Don't route signal traces across a split or across an unrelated plane.

Footnotes:

1. Johnson, Howard, and Graham, Martin; High-Speed Digital Design, A Handbook of Black Magic; Prentice Hall, 1993, p.191
2. See also "Ground Plane 101", October, 1997, p34.
3. "PCB Impedance Control, Formulas and Resources", March, 1998, p12; Note that the impedance calculations always include the distance between a trace and its reference plane as one of the variables. But in the special case of differential signal lines, differential impedance is sometimes achieved without the use of planes.
4. "Loop Areas, Close 'Em Tight", January, 1999, p. 22
5. "Crosstalk, Part 2: How Loud Is It?", December, 1997, p.54. See also Howard Johnson, op. cit. Figure 5.4 at page 192.
6. "ESR and Bypass Caps, When Less Might Be Better", June, 2000, p. 30.
7. Rick Hartley, "Controlling Radiated EMI Through PCB Stackup", July, 2000, p. 16.
8. See Analog Devices "High Speed Design Techniques" Section 7b:"Grounding in High Speed Systems", p 2 at: http://www.analog.com/support/standard_linear/seminar_material/highspeed/highspeed.html
9. See for example Analog Devices Application Note "CMOS 240 MHz Triple 10-Bit High Speed Video DAC ADV 7123", p. 16
10. For a related discussion, see "Slots In Planes, Don't Use 'Em", March 1999, p. 36. See also Howard Johnson, op. cit. p. 194ff.