



PCB Layout Considerations for Switchers

Chris Richardson

Applications Engineer

**Medium and High Voltage Power
Management**



Layout Considerations Topics

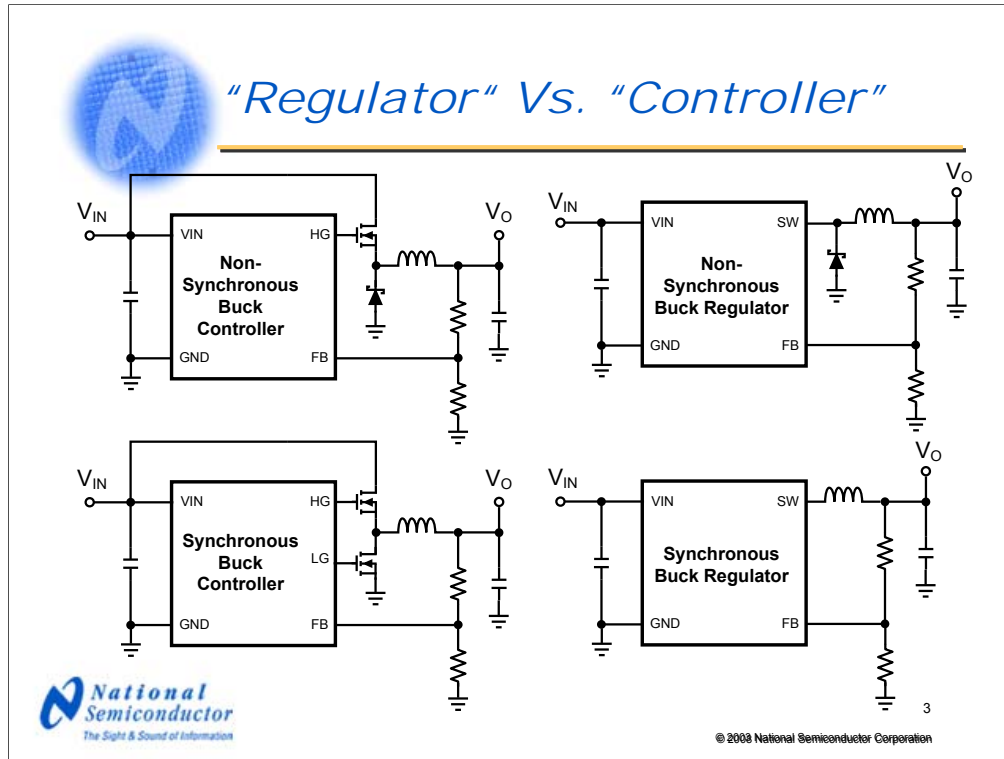
- Self inductance influences impedance
- Noise Coupling Mechanisms
- Locating the high di/dt loops
- “Ground rules”
- Copper Requirements For High Current Paths
- Component Placement Strategy
- Gate Drive Layout Requirements
- Power FETs and Decoupling
- Switch Node design
- Output Capacitors
- Control Circuit Consideration
- Noise considerations
- Thermal Considerations



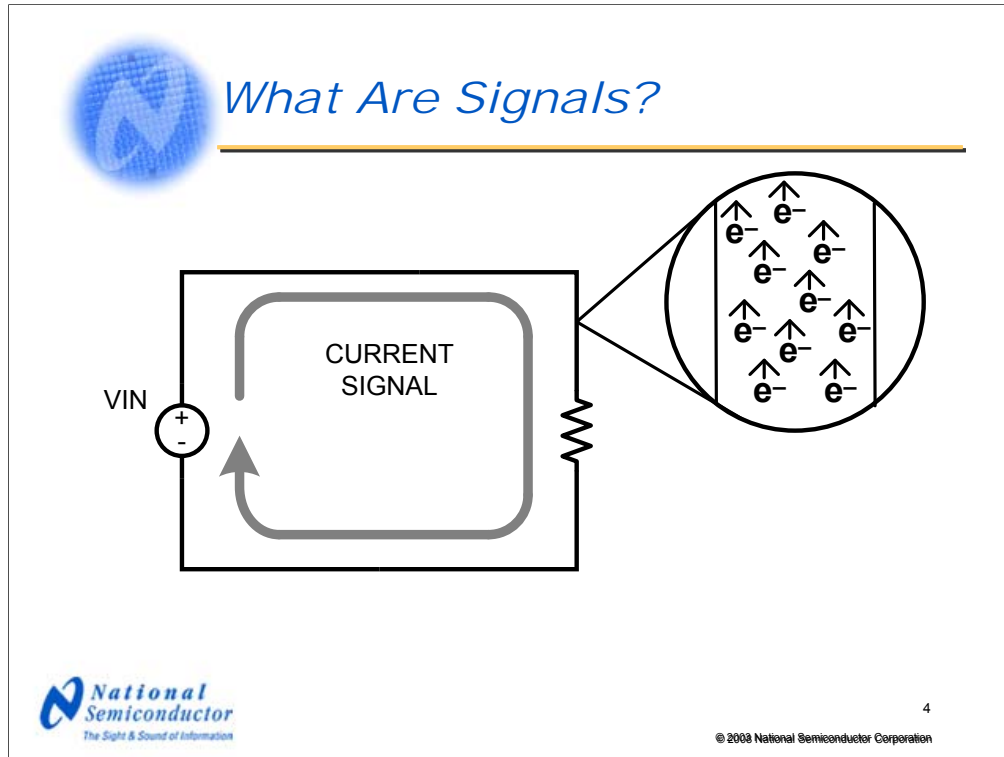
2

© 2008 National Semiconductor Corporation

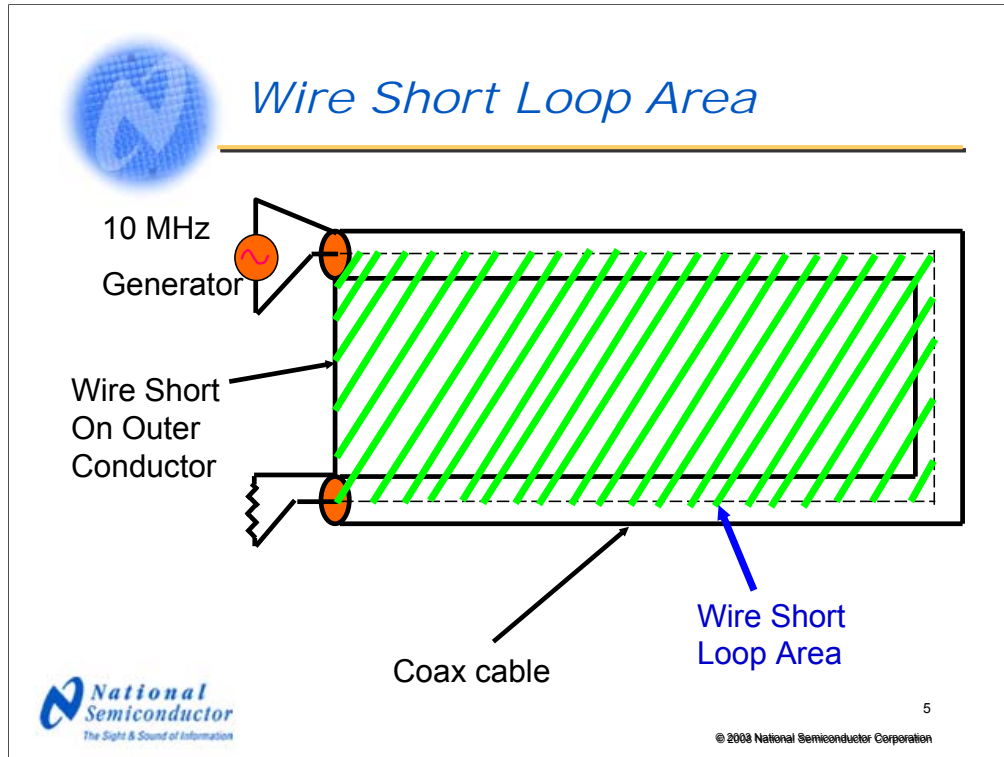
The following is a list of the topics to be covered. We'll discuss noise sources and coupling mechanisms and then see how proper PCB layout can be used to minimize these effects. Before we begin I want to explain a few terms that factor heavily in this presentation. 'dV/dt' refers to the rate of change of voltage, and 'di/dt' refers to the rate of change of current. The square wave and trapezoid wave currents and voltages encountered in hard-switched converters such as the classic buck, boost and buck-boost regulators are good examples of circuits that have high dv/dt and high di/dt. Although switching regulators operate at relatively low frequencies (typically between 100kHz and 3MHz) the very high di/dt and dv/dt signals have frequency components that may reach upwards of 100MHz or more. As such, we must treat a switcher as though it is truly a high frequency RF circuit. What's worse, is that in many cases there is significant energy contained in these high frequency components. Treat them with respect or you will be sorry!



Also, within National Semiconductor, when we talk about a switching regulator we are referring to a device with at least one internal power switch. A switching controller, on the other hand, is a device that controls external power switches. In both cases the power switches are nearly always MOSFETs. When we refer to 'synchronous' regulators we are talking about switching converters that use a MOSFET instead of a Schottky diode for the second, or 'synchronous' switch.



I was once fortunate enough to attend a training by a colleague of mine who had devised some clever visual demonstrations of noise measuring and noise removal techniques. One thing he said during that training that really struck me was that signals are not voltages traveling over wires. Signals are currents that move from a source of electric potential, and must return to that source in order to flow. It helps me to visualize actual electrons and holes moving in the copper planes and through the components of my circuits when I am designing my PCBs.



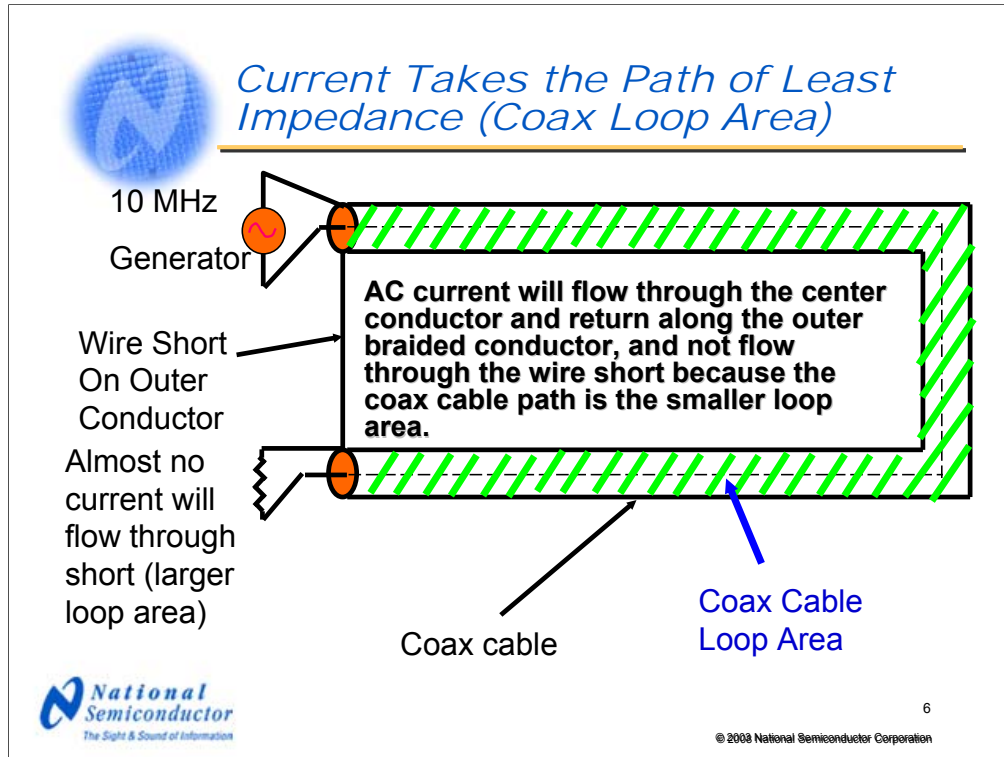
Current takes the path of least IMPEDANCE, NOT of least RESISTANCE. At DC, the paths of least resistance and least impedance are equal, but at any frequency greater than zero they are different. In practice frequencies of about 3 kHz or more exhibit this difference. Remember these points:

The path of least impedance will not necessarily be the path of least reactance.

The path of least impedance will be the path with least self-inductance.

The path with least self-inductance will be the path with smallest loop area.

A signal generator connected to a piece of coax cable with a short circuit between the braids formed the basis of a visual aid that made this concept clear to me.



Two current probes were set up to measure the currents flowing out of the signal generator and flowing through the wire short. That wire short was actually a piece of solder wick saturated with solder, which by the way makes a great zero ohm resistor in a pinch. At DC the two currents were nearly equal, but as the frequency increased less and less current went through the wire short because of the large loop area, which had more inductance. Comparing the area made by the wire short and the center conductor with the area of the coax outer conductor and the same center conductor, the loop area of the coax was much smaller than the loop area with the wire short. At frequencies above a few kHz the return current path was through the coax outer conductor because it was the path of least loop area/least inductance/least impedance, and it was clearly visible on the oscilloscope. That made a believer out of me.

Self Inductance

CURRENT IN →

Inductance, L

CURRENT OUT ←

d

S

$2a$


$L = \frac{\text{Magnetic Flux}}{\text{Current}}$

$L = \frac{\mu d}{\pi} \left[\ln \frac{S}{a} \right]$

Assumes $d \gg S \gg a$

L gets BIGGER as loop area increases

L gets SMALLER as wire diameter increases

 National Semiconductor
The Sight & Sound of Information

7

© 2008 National Semiconductor Corporation

PCB traces have inductance. Thicker traces (effectively bigger wire diameter) will reduce the inductance, but only slightly. Remember that for current to flow in a circuit there must be a loop, and the area enclosed by the current loop has the greatest effect on stray inductance. Keeping current paths directly adjacent to their return paths, parallel to one another, or on the same layer or directly above/below on adjacent layers is an effective way to reduce stray inductance.

If this presentation had a subtitle, a good choice would be "The Inductor: Your Friend in it's place, your Worst Enemy Elsewhere."



How Does Noise Couple in a System?

There are four and ONLY four mechanisms for noise to propagate through a system

- 1) Conductive**
- 2) Near field magnetic (transformer)**
- 3) Electric field (capacitor)**
- 4) Near and Far field electromagnetic (radio)**



8

© 2008 National Semiconductor Corporation

Before we discuss PCB design techniques to minimize noise issues, let's review the sources of noise, listed here in order of how likely they are to be present in YOUR system. Conducted noise is by far the most common enemy. It generally is produced when high switching currents encounter stray inductance in the power path, producing large voltage spikes and high frequency ringing. Anytime two circuits share a common conductor (such as a ground plane) they have the potential for interaction between them. Near field magnetic effects are caused by direct flux linkages between magnetic fields and conductors. Any conductor with high switched current will have an associated magnetic field and therefore has potential to couple the signal (or noise) it is carrying to an adjacent conductor. Electric fields couple through capacitive effects. The coupling is proportional to the area of the source and victim circuits that are in close proximity. The higher the dv/dt of the signals being carried by the source of the noise, the better the coupling efficiency. The thing that's probably blamed the most but is in fact the least responsible for circuit operation problems is radiated fields. The higher the di/dt in a conductor, and the larger its loop area the better transmitter it becomes. Due to the extremely long wavelengths (relative to the size of most circuits) of low frequency signals, radiated interference is only likely for very fast moving signals. Minimize loop areas to minimize the quality of the unintended antennas you are creating.



Conductive Coupling

- **Common Impedance**
- **Requires two or more conductive contacts**
- **This mode is responsible for over 90% of noise problems!**

Examples

- **AC or DC power leads**
- **Grounding connections**
- **A shared signal path**



9

© 2008 National Semiconductor Corporation

PCB traces, like wires, have resistance, self-inductance, mutual-inductance and capacitance to adjacent traces. A common problem with ground noise is having a noisy node directly connected to a ground plane. Your ground plane is useful as a reference only if there is precisely zero current flowing in the ground plane. If there is any current at all flowing through the ground plane, there will be a voltage gradient across the plane. Two circuits now connected at different points along the ground (reference) plane will see a different reference voltage and the potential for cross talk exists. In practice we allow DC currents to flow in a typical ground plane, but keep noisy current signals out of the system ground plane if you want to avoid noise problems!

**Near-Field Coupling -
Electric (E) or Magnetic (H) ?**

Line Picking Up Noise

**V Causes E
I Causes H**

Guard trace: Tie to ground, power, op amp output. Whatever is appropriate for a given situation.

Disconnect the LOAD!

If no current is flowing and the noise problem is still present then the coupling mode is Electric Field (E)

National Semiconductor
The Sight & Sound of Information


10
© 2008 National Semiconductor Corporation

There's a simple test to find out if near-field coupling is being induced by electric fields or magnetic fields. Remove the load, or reduce the output current to a minimum. Magnetic fields cannot exist without flowing current, so if the noise diminishes or disappears, it must be due to a magnetic field. If on the other hand the noise stays the same at no-load, the source is an electric field. To reduce electric-field coupling place a guard trace between the radiating trace and the trace picking up noise. Tie the guard trace to a low impedance with the same voltage potential as the circuit picking up the noise. We tend to think of ground as being the only low impedance point, but input and output voltages are also low-impedance for AC thanks to the large amounts of filter capacitors.

Reducing magnetic-field coupling is not easy to achieve with shielding. Instead, stop the problem at the source by minimizing loop area, giving preference to the heaviest currents that are switched. Separate the distance between circuits, use self-shielding techniques to allow the return-current conductor to surround the signal conductor. Think about coaxial cables, which completely enclose their send path with the return path. Toroidal or pot Core inductors are also helpful tools for reducing stray magnetic flux.

Magnetic Self Shielding

- **If the return current can completely enclose the signal current, excellent magnetic field containment will result**
- **Any techniques that will reduce self inductance will also reduce or eliminate magnetic fields**

 National Semiconductor
The Sight & Sound of Information


11
© 2008 National Semiconductor Corporation

Magnetic flux is inversely proportional to the distance from the conductor carrying current. Conductors carrying equal and opposing currents will create fluxes that cancel each other. By enclosing a signal with its return path (like in a coax cable) the flux is practically enclosed in the wire itself and will not disturb neighboring circuits. Note that each of the conductors produces precisely the same magnetic flux that it would if it were separated from its return conductor. The secret to success is the fact that the fields in the source and return conductors are equal and opposite. When placed in close proximity, they tend to cancel. Always try to optimize field cancellation in your circuit layouts. If we could, we'd build our circuits from coaxial cable, but this isn't practical. In the mostly two-dimensional world of PCBs the idea is to keep signal and return paths close to one another, and this can be side-by-side on a single layer or one above the other on two adjacent layers.

*How Hidden Antennas
Are Made on a PCB*

Two Layer Board

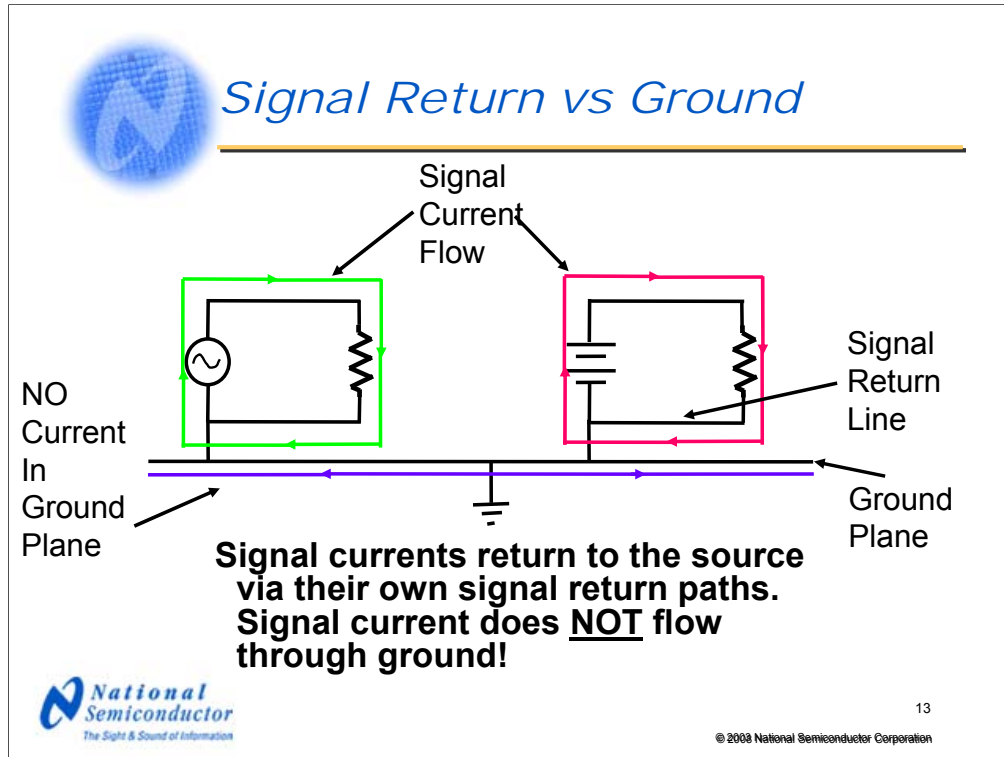
Row of vias opens up ground plane between pins and causes AC- return current to flow around gap in plane. The gap will radiate at whatever frequency is $\frac{1}{4} \lambda = \text{gap length}$!


12
© 2008 National Semiconductor Corporation

Imagine that this diagram represents a switching regulator, labeled 'OSC', and a digital load like an FPGA or DSP, labeled 'IC' since these are the most common type of load for our switchers. The black arrows represent the actual top-layer copper traces that the current takes from the switcher to the load. That current has both a DC portion and an AC portion, a fact of life for all switching power supplies. The load consumes this current – which in reality means that the current passes through the load on it's way back to the source. It's at this point that the AC and DC parts may well diverge. A solid plane of copper connected to ground fills the bottom layer of the PCB, so the red line represents the return path of the DC current. If we could somehow get a current probe into the plane this would be plain to see. The green arrows show where the AC portion of the current returns, and the higher the frequency, the more closely those electrons will follow the traces on the top layer as they follow the path of least inductance. Trouble can ensue if a via interrupts the AC current flow. Keep an eye on the path of least impedance!


About far field (electromagnetic) coupling. To minimize electromagnetic radiation or pick up:

1. Keep antenna surfaces free of common mode current.
2. Minimize antenna area.
3. Keep antenna length $< \lambda / 20$.

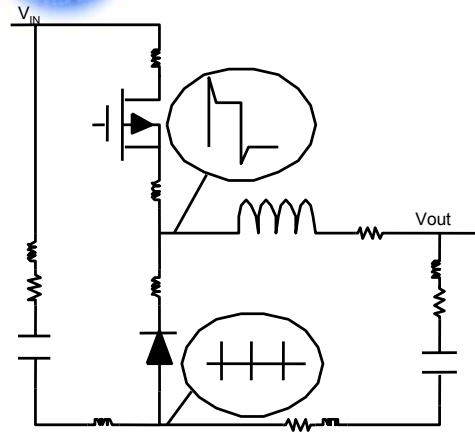


The basic rule of thumb to employ if you have currents flowing in your ground plane is:
Ground isn't!

Remember, all conductors have resistance, even planes! Signal currents flowing through a ground plane will cause a voltage gradient across the ground. To ensure the ground reference is the same across the board, it is better to have signal return paths kept locally, as shown. Imagine what the ground currents would look like if the two resistors depicted above were simply connected to the ground plane with a simple via dropped through a pad.




Impact of High di/dt



- All elements, including PCB traces, have parasitic L, R, C
- High di/dt thru parasitic L produces voltage spikes
- Must avoid injecting these currents into the ground plane

First, identify high di/dt paths



14

© 2008 National Semiconductor Corporation


If you think about the buck topology, you have to ask the question: Why should there be spikes on the output and where do they come from? There is an inductor, generally operating in continuous-conduction mode, in series with the output. So if signals are currents, and spikes are very fast-moving signals, then by definition there are no high di/dt currents allowed to flow through to the output capacitor. In reality, the observed spikes are frequently very large. How can this be, and how do we control the problem? The source of the problem is usually high di/dt currents being injected into the ground plane and polluting the entire circuit board; EMI wise that is.

High di/dt currents will cause disturbances on our boards that show up as spikes all over the circuit. They tend to have very fast edges and therefore very-high-frequency components associated with them. (Take a moment to ----- remember Fourier analysis). High frequencies tend to radiate better than low frequencies and will seem to be present wherever we try to make a measurement, especially if they get injected into the ground plane. So the goal is to keep the ground plane quiet!

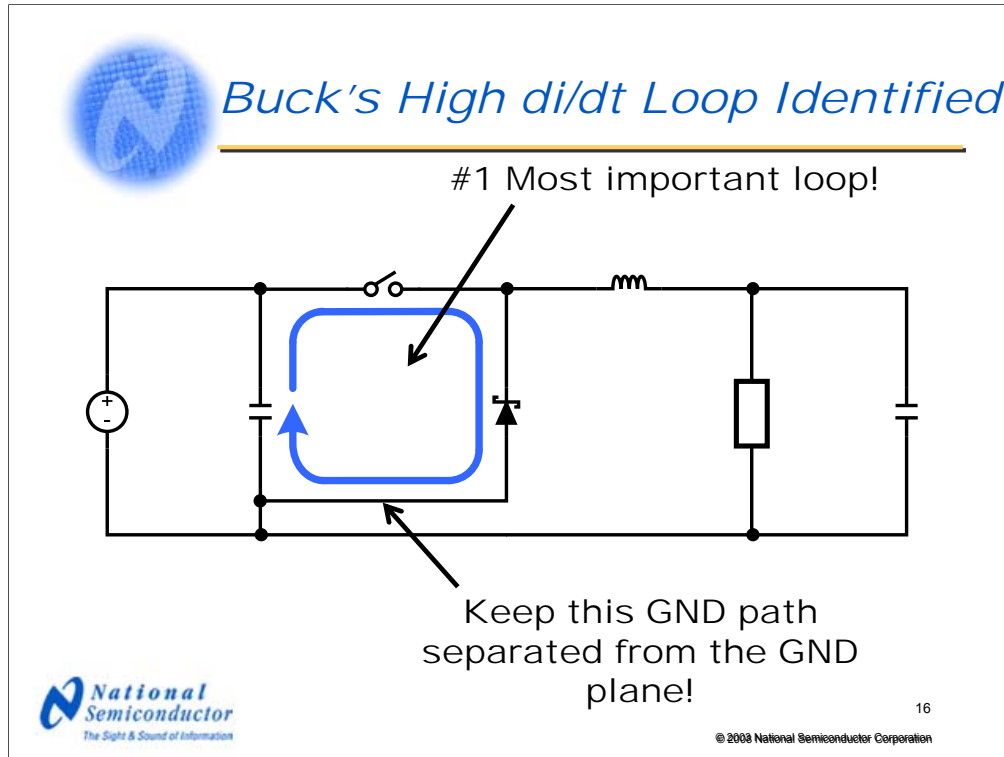
-----The solution is to control the path taken by high di/dt currents. But before we can decide how to control these trouble makers, we had better have a methodology for determining where they are in the first place.

High di/dt Loops in a Buck

- Draw the switch-ON current path in one color
- Then draw the switch-OFF path in another
- Any part of the circuit that has only a single color is a high di/dt path
- Works for all topologies


15
© 2008 National Semiconductor Corporation

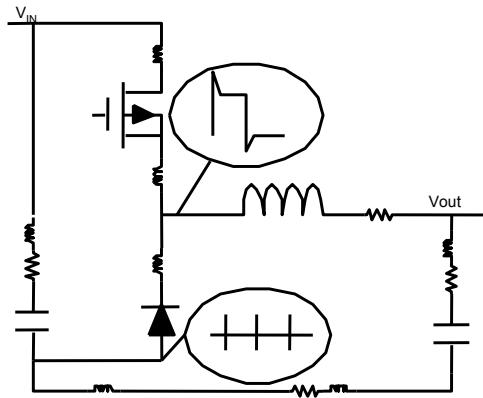
The diagram above shows a neat trick to help find the heavy switching currents. Simply look at the current paths in the various switch states and draw the loops in a different color for each switch state. For instance, in the buck example above, with the switch on, current flows from the input capacitor, through the switch, the inductor, and then the output capacitor and back through the ground plane to the input cap return. On the second half of the cycle, the inductor is forcing current to flow through the output cap, ground, the catch diode, then back to the inductor. Draw the first loop in red, for example, and the second in green. Now look to see where only one color exists on a path. The single color paths are the high di/dt paths, aka the heavy switched currents. They will need your undivided attention during the layout phase!



The blue arrow indicates the key part of the buck regulator layout. This is the path which takes priority over all others, and this is the loop which should be made the smallest. Another way to say this is that the inductance in this path should be minimized before any other loop is laid out. First, the input cap needs to be as close to the switch and diode as possible. Second, the anode of that diode (or source if it's a MOSFET) needs to connect to the input cap before it is tied to the GND plane.

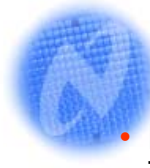


How to Deal with the Noise Generators



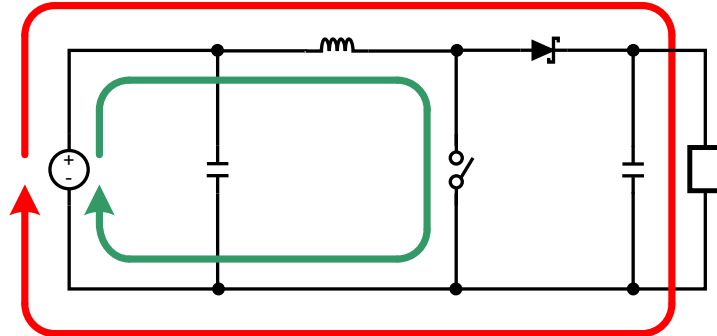
- Note the re-route of the diode-return path
 - Forces pulse currents directly back to input cap
 - Keeps high di/dt currents out of ground
 - Diode anode may actually be a bit noisier, but who cares?
- Can apply the same rationale to all topologies

The fix in the buck example is to route the diode anode directly back to the input cap ground terminal. This forces the ground current to flow in a continuous manner that is in fact equal to the inductor current. If done properly, the path length of the anode connection is not significantly longer than if it were simply connected to the ground plane. Of course in a synchronous regulator all the above reasoning applies to the low-side-FET source connection. For a boost topology, simply imagine swapping the FET and the diode in the circuit above and also swapping the input and output labels. The buck has been turned in to a boost and the current paths look essentially the same as they did before. So the proper place to return the switch ground is the bottom of the output cap as opposed to the input cap for the buck.

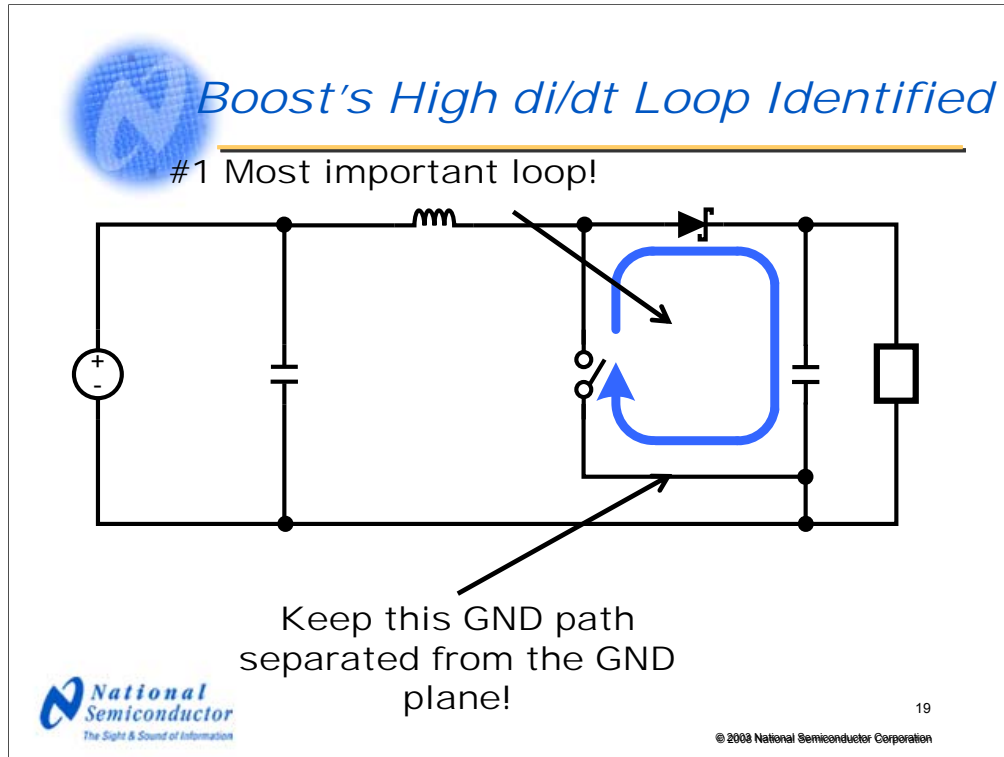


High di/dt Loops in a Boost

- Draw the switch-ON current path in one color
- Then draw the switch-OFF path in another
- Any part of the circuit that has only a single color is a high di/dt path
- Works for all topologies



The exercise is repeated here for the boost converter, where the heavy switching currents are at the output of the converter.



Again the blue arrow indicates those parts of the circuit that carry heavy current during one portion of the switching cycle. Now, with a boost regulator, to keep the ground plane free of electrical noise the source of the low-side FET should connect to the negative of the output cap before it is electrically connected to the GND plane.



High-Current Copper Requirements 1

- **Do not use minimum width traces**
- **Approximate trace width as follows:**
Where T = trace width in mils, A is current in amps, and CuWt is copper weight in oz

$$T = \left(-1.31 + 5.813 \times A + 1.548 \times A^2 - .052 \times A^3 \right) \times \frac{2}{\text{CuWt}}$$

- **Formula works over a range of 1A to 20A**



20

© 2008 National Semiconductor Corporation

Most designers tend to ignore the trace width since it's generally not a concern for digital applications. The usual approach is to try and design for minimum trace width. This will get you in lots of trouble with high current designs. The formula above is derived by curve fitting the standard mil-spec curve that's been published for decades now. It's a reasonable estimate over a fairly wide range of currents. In general, where you're looking at high current paths, more copper is better. When I'm designing my PCBs for switchers I don't spend time calculating this formula, I connect almost everything with copper pours or "shapes".

There's a tendency to use ½ oz copper on top-side high-speed layers since you can etch to finer pitch with the thinner material. But that's not a good power strategy. Where possible, at least try to get 2 oz for inner plane layers where fine pitch is not an issue. Most PCB fab houses are able to do selective plating on outer layers. This can be done to thicken up the copper in the high-current paths but adds cost and isn't done very often. One way around this is to use multiple layers to carry heavy currents and then connect those layers with lots of vias in parallel.



High-Current Copper Requirements 2

- **Some examples:**
 - 1A, 1 Oz Cu, trace width = 12 mils min**
 - 5A, ½ Oz Cu, trace width = 240 mils min**
 - 20A, ½ Oz Cu, trace width = 1275 mils min**
- **Clearly lots of width required for high currents with light weight copper planes**
- **These widths are designed for an approximate 10°C temp rise. Wider is better!**
- **Try to design for 30 mils per amp for 1 Oz Cu and 60 mils per amp for ½ Oz Cu**
- **Shapes with switching currents should be wider**



21

© 2008 National Semiconductor Corporation

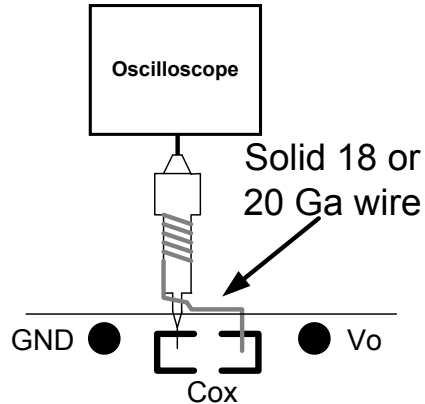
A 50mil, ½ oz trace may fuse open when used at more than 10A. Keep in mind too, that the top side copper in particular is the primary thermal path to the cooling air flowing over the PCB. Since the thermal resistance of copper is much lower than the fiberglass board material, the more copper that's left in place, the lower the overall thermal resistance of the assembled board.

A former manager of mine who helped write much of this presentation used to review all of our PCB designs, and one of the things he would do was to hold a PCB up to the light to see how solid all the planes were. He liked to say that the price is the same regardless of how much of the board is covered by copper, so get your money's worth and extend those power and ground planes.



Component Placement Strategy 1

- Design power path layout first. Keep multi-phase layouts symmetrical
- May help to work from the output back towards the Input source, or starting with the switch node
- Locate sense resistors and Inductors, then FETs and input caps.
- **Keep copper width requirements in mind!**



Take outputs right across final output cap



22

© 2008 National Semiconductor Corporation

There's a tendency to think that the controller is the most important part of the design and so that's the place to start. Nothing could be further from the truth. You want to start with the power path components since they are large and require fat interconnects. It's a lot easier to find ways to sneak in little signal traces than the large plane areas associated with the power path.

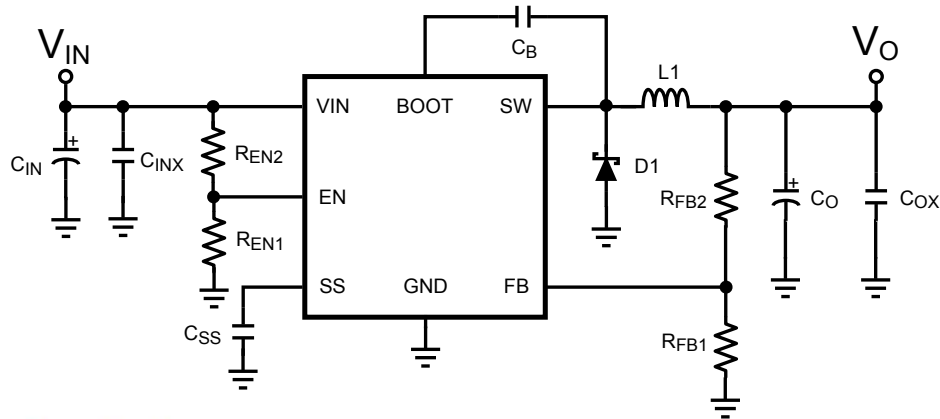
Lay out the power parts in such a way that the power flow makes logical sense and the enclosed loops are as small as possible. Try and force return currents to flow either next to or under their respective source current. This minimizes the loop area and reduces the magnetic fields that will radiate from the board's surface. The equal and opposite fields produced tend to cancel one another. The closer the conductors, the more effective the cancellation.

Once all the power components have been placed, look for quiet areas to locate the control circuits.



Step-by-Step Layout

- **Non-synchronous buck regulator**

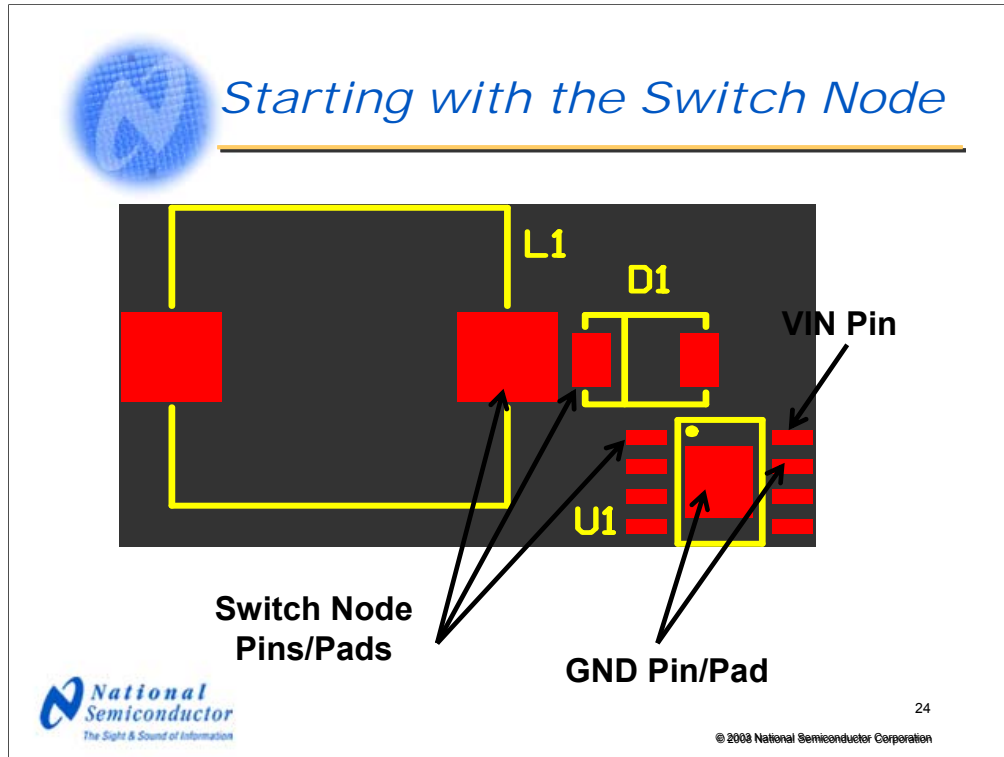


National Semiconductor
The Sight & Sound of Information

23

© 2008 National Semiconductor Corporation

Circuit schematics like this example rarely provide any insight into proper PCB layout, although I have seen some footnotes placed near certain components on various customer designs. Like many schematics, the pinout shown is not representative of the actual IC, which is a power SO-8 with a thermal tab underneath.



The switch node is a great electrostatic radiator as it swings between VIN and ground at the switching frequency. Displacement currents will flow in an effort to charge and discharge the capacitance between the switch node and the ground planes and free space. In general, it is better if the area of this node is minimized. I like to start with the power switch, diode, and inductor that make up the switch node and place them so that their pads are as close as possible, and then I put a shape that just covers all three pads. For other than buck regulators, take a look at any nodes that swing through high potentials and look for ways to minimize the capacitance to the rest of the world.



Power FETs and Decoupling

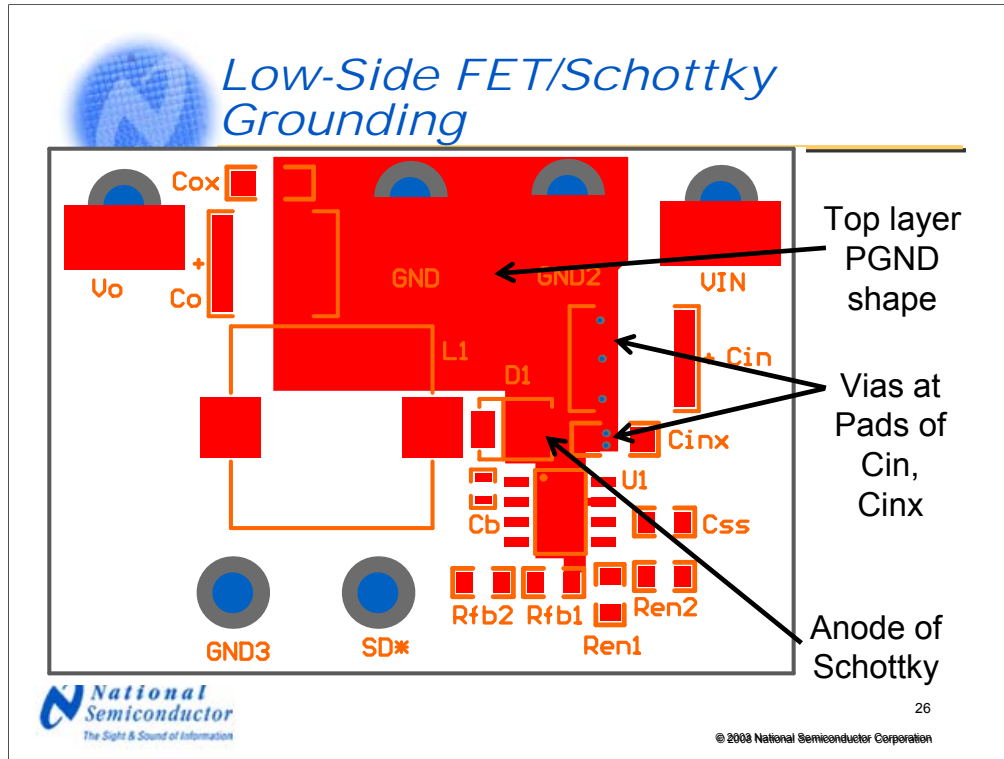
- **Minimize loop area enclosed by high-side FETs, low-side FETs, and input caps**
- **Connect the low-side FET's source to the input- cap ground, then to the ground plane**
- **Use copper pours for drain and source connections to power FETs**
- **Use lots of vias to tie in to inner layers**
- **Minimize stray Inductance in the power path!!!!**



25

© 2008 National Semiconductor Corporation

As discussed earlier, the loop between the input cap, high-side FET and low-side FET needs to be minimized. This is probably the most critical path in buck designs. Follow the guidance discussed about low-side FET-source grounding at the input cap ground to minimize output spikes and be sure to use large copper pours for the main power connections to the MOSFETs. This is helpful not only electrically, but also thermally. The same philosophy applies to the use of vias. More is better. Just think of the inductances all connected in parallel. They also act as thermal passages to the back side and inner layers of the board. Fill with solder if possible. While not nearly as good as copper either electrically or thermally, solder is much better than air and that's what would fill the holes otherwise.



A word of caution to emphasize the point about the low-side FET or the anode of the diode. It is a noisy node and should not be directly connected to the ground plane. Doing so will inject noise into the ground plane and corrupt it for the rest of the nodes depending on it. If you see spikes on the output of a buck regulator, this connection was probably made incorrectly. In an ideal design there would be nearly no spikes on the output of a buck regulator. In general, any spikes that do appear on the output are conducted in through the ground rail. With careful attention to the path between the low-side switch grounding to the input capacitor, output spikes can be nearly eliminated or at least greatly reduced from what is commonly accepted for a switching regulator.

Connecting Bypass Capacitors 1

- **Connecting to output bypass caps**

Terrible!

Good

National Semiconductor
The Sight & Sound of Information

27

© 2008 National Semiconductor Corporation

Take a look at the two layouts above. An output capacitor is intended to be connected to the surface mount pads in each case. The design on the left will not do a very good job of reducing ripple and spikes due to the high inductance connections to it. The straight through connection paths from the input to the output is likely a lower impedance path than the path through the capacitor. In the case on the right side, the input and output traces are much heavier so they will be lower impedance connections. But the more significant consideration is the physical orientation of the traces relative to the capacitor. Note that the current path is through the capacitor pad and the output connections come directly from the capacitor pads as well. This is the lowest impedance connection viewed from the perspective of the output looking back into the capacitor. The next improvement that could be made would be to make the traces wide enough that they effectively become a couple of planes. If multiple decoupling caps are used, it's best to create top-side shapes that all the parts sit on. If the capacitors need to connect to internal ground or power planes, use a large number of vias to connect these shapes to the inner layers.



Connecting Bypass Capacitors 2

Connecting to high frequency bypass caps: This assumes a connection into internal planes

- | | | |
|--|------------------|--|
| | Terrible! | <ul style="list-style-type: none"> • Long thin traces add inductance and effectively isolate the capacitor |
| | Good | <ul style="list-style-type: none"> • Keeping vias close to pads minimizes parasitic inductance |
| | Better | <ul style="list-style-type: none"> • Doubled vias further reduce inductance |
| | Best | <ul style="list-style-type: none"> • This technique further reduces inductance by reducing the high frequency loop area |



28

© 2008 National Semiconductor Corporation

High frequency bypass capacitors need to be connected with minimum added inductance, but there's a conflict here between the optimal layout and the manufacturability. Many times design rules in place tend to force solutions like the top design shown above. That capacitor will be a complete waste of money and space since the added inductance will effectively decouple it from the load. If possible we would put two or more vias right in the center of the capacitor pads, but this can cause problems with reflow soldering when the vias wick away the solder. The best practical approach is shown on the bottom of the slide. The enclosed loop area is minimized and there's good field cancellation. This makes for a low inductance interconnect.



Component Placement Strategy

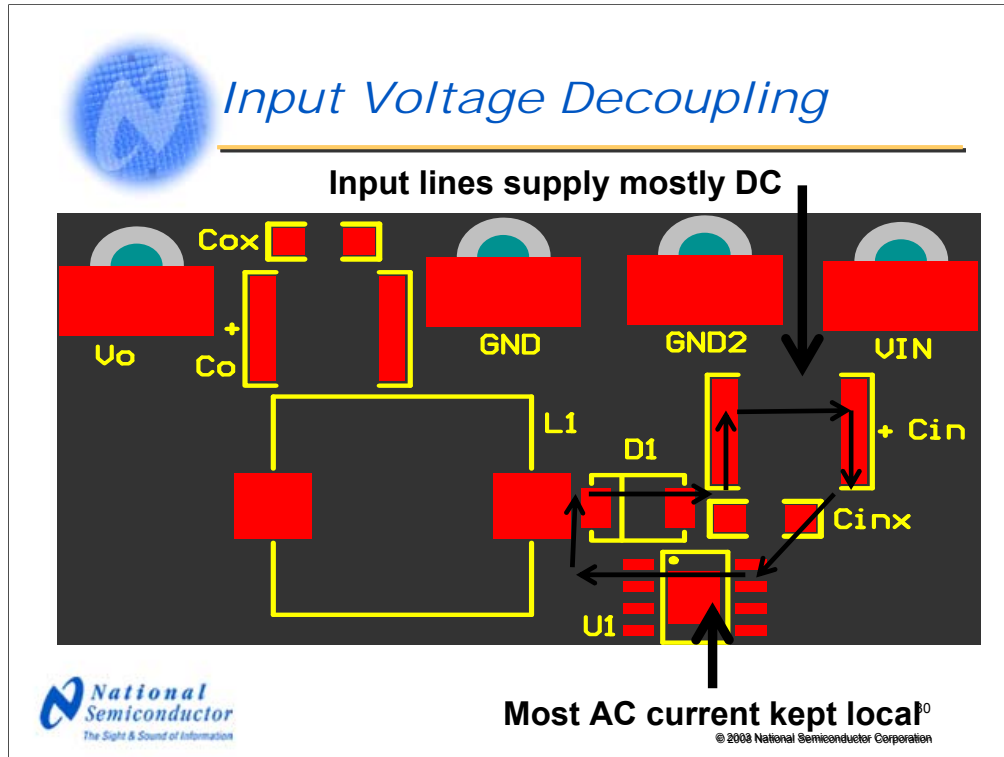
- **Control circuits go in last. Small traces, easy to route**
- **The higher the impedance and/or gain, the smaller the node should be, especially FB pins, input to op-amps, comp pin, etc**
- **Low Z nodes can be big, including outputs, so put FB components near the inputs**
- **Play with some “Paper Dolls” to see the power flow**



29

© 2008 National Semiconductor Corporation

You may have seen small RC filters at the input to the VCC pins of some switching regulator controllers and thought that the idea was to prevent the controller from creating noise on the VCC supply. In fact the opposite is the case. In general such a supply will be used for gate-drive power as well as possibly logic in the main system. This tends to be a very noisy rail and so the RC filter is used to filter the supply voltage to the controller! Therefore, keep the VCC bypass caps as close as possible to the controller's VCC pin and connect the other end directly to ground, preferably at the controller's ground pin. Whether the bypass cap should connect to AGND or PGND will depend on whether the VCC rail supplies only control circuits in the IC or gate drives. If the former, connect to AGND. If the latter, go to PGND.



Probably the single biggest cause of radiated emission problems from switching power supplies is conducted emissions on long input power lines. Shielding can help, but the ultimate fix is to reduce the AC currents in the input lines. This implies good decoupling and filtering of the input currents to the regulator. Keeping the loop between input caps and FETs as small as possible is another key way to keep this under control. Use the capacitor connection techniques discussed in the previous slides for output caps. They work at the input side as well! It's worth creating a set of "Paper Dolls" for your main power components and lay them out to see how the main power path interconnects will flow. This can save hours of grief later when things don't fit the way you'd like.



Control Circuit Layout

- **Make long runs to low-impedance nodes, short runs to high-impedance nodes**
- **Route current-sense traces parallel to one another – minimize differential-mode noise pickup**
- **Keep most small signal traces thin – lower capacitance to surrounding signals**
- **Route Sense+/Sense- as a parallel run**



31

© 2008 National Semiconductor Corporation

The lower the AC impedance of a signal run, the less likely it is to be corrupted by stray fields. Therefore, if given a choice as to which end of a component to make the long run, always choose the lowest impedance end. An example is the feedback divider top resistor. Given a choice of locating the resistor close to the high impedance error amp input and the low impedance output connection, always place it close to the error amp and make the long run to the output voltage. When running current-sense traces, try to keep the two lines parallel to one another and as close together as possible. That way, any noise picked up will likely get picked up equally and cancel differentially. It will look like a common-mode signal and so will be subject to the CMRR of the receiving amplifier. But at least you won't be injecting purely differential noise signals. The same rationale applies to such things as remote sense lines. Keep the small signal traces small. They will be less susceptible to capacitively coupled noise pickup. In cases where long runs must be made with high-Z signal paths, place the conductors parallel to one another on an inner layer and surround with grounded copper on all sides. In essence, build a PCB-coax cable.



Control Circuit Layout

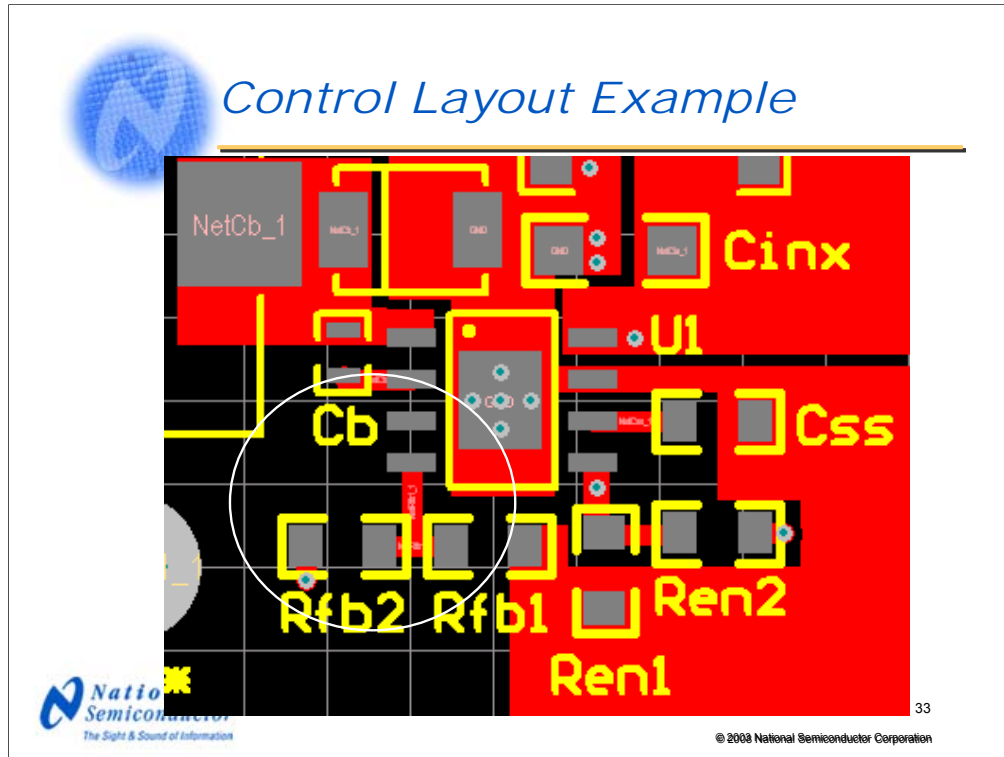
- **Use single point ground for AGND-PGND connection**
- **Can use top side “daisy chained” ground or separate plane area for Sgnd connections**
- **Keep V_{CC} bypass cap close to pins**
- **Route sensitive signals away from noisy nodes, and no noisy signals near sensitive nodes**
- **Keep feedback Rs and Cs close to pins**



32

© 2008 National Semiconductor Corporation

It's generally desirable to tie the controller IC's AGND and PGND pins together at a single point that also ties to the ground plane. If you have room for a dedicated, top-side ground shape, you can use that as the return for all ground referenced control circuits (feedback divider, softstart cap etc.). In most cases that will not be a viable option. Using a small trace that daisy chains from the AGND pin to the various small signal grounds is an excellent alternative. If you need to route this trace through multiple layers there will be a problem with the layout software wanting to tie the via to an internal ground plane. That would destroy the whole point of wanting to segregate this trace. You need to find a way to fool your software into doing what you want. One approach is to come up with a separate AGND symbol. This allows you to tie all the signal-ground points together thru multiple layers without connecting to the internal plane. However, you will generate a DRC error when you try and connect the AGND and PGND. In some programs you simply accept the error and ignore it knowing full well what it is. Some programs will not allow you to connect the two separate nets on the schematic. If you make a connection between the two nets with a "line" as opposed to a "wire" the program will allow the nets to coexist. The appearance of the schematic is good, but you will generate a DRC error on the PCB.



Follow the grounding suggestions discussed earlier. Keep the highest impedance traces short since they will have the greatest tendency to pick up stray fields. So as an example, the two resistors in a feedback divider should be located very close to the FB pin of the regulator, not near the power supply's output. The connection from the top of the divider to the supply output is a very low impedance connection, and as such will be quite immune to noise pickup. The error amplifier input, in contrast, is a very high Z input and will be very susceptible to stray field induced noise. There's a tendency to want to make sensitive nodes large in the mistaken belief that this will offer some shielding effects. In fact the opposite is true. This increases the capacitance to space and increases the likelihood of noise pickup. Make sensitive traces narrow and as short as possible. These are where you want those 10mil-wide traces.



Gate-Drive Layout

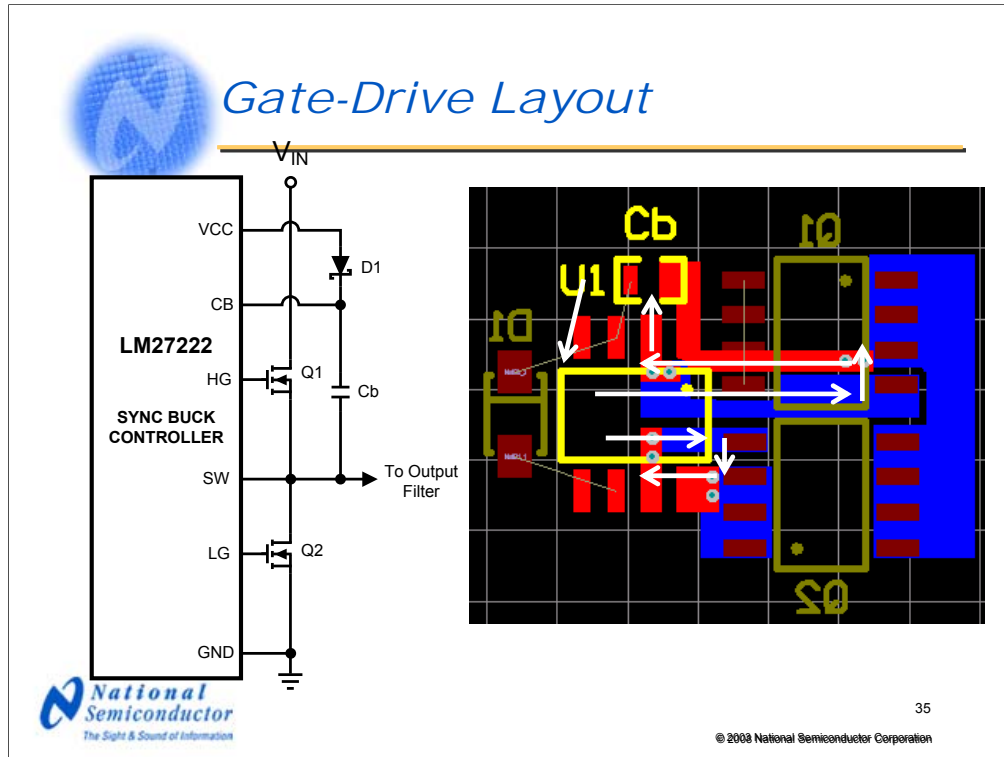
- **Place drivers close to MOSFETs**
- **Keep C_{BOOT} and V_{DD} bypass caps very close to driver**
- **Minimize loop area between gate drive and its return path: low inductance**
- **SW-pin connection should be 0.015" wide or larger, as should the gate connections**
- **Minimize stray inductance in the power path!!!!**



34

© 2008 National Semiconductor Corporation

The gate connections to external FETs are very high di/dt paths and as such need careful thought. Worse yet, the connection between driver ground and FET source will act like a parasitic multiplier since the inductance will act as a gain-degeneration element for the power device. Long connections here will dramatically slow the rise/fall times of the FET switches. This is particularly important for the high-side FET in buck applications. Try and lay the gate-to-driver connection directly over the source-to-driver ground connection to minimize the area enclosed by the resulting loop and use traces at least 15 mils wide (0.3mm).



In the partial layout on the right the components have been spread out more than necessary for clarity. In an actual design the FET driver, boot cap and boot diode could be almost directly on top of each other using a two-sided layout. In a synchronous buck converter like this one, the high-side FET, Q1, suffers lots of switching losses, whereas the low-side FET, Q2, suffers almost no switching loss. The emphasis should be placed on minimizing that high-side FET loop from the boot cap Cb to the gate pin of Q1 and back.

When paralleling FETs for very high current applications, keep the loops between the driver and FET gates equal in size, so that each FET turns on and off at the same time.



Component Placement Strategy – Rules Of Thumb

- You can put inductance in series with an inductor, but do it on the quiet side, NOT the switch node side
- 30 mils per amp for 1 Oz Cu and 60 mils per amp for ½ Oz Cu
- 1 A of DC MAX per via is a good design goal.
- Vias to bypass caps should be placed tangent to the pad, two per pad is preferred
- **Minimize Stray Inductance in the Power Path!!!!**



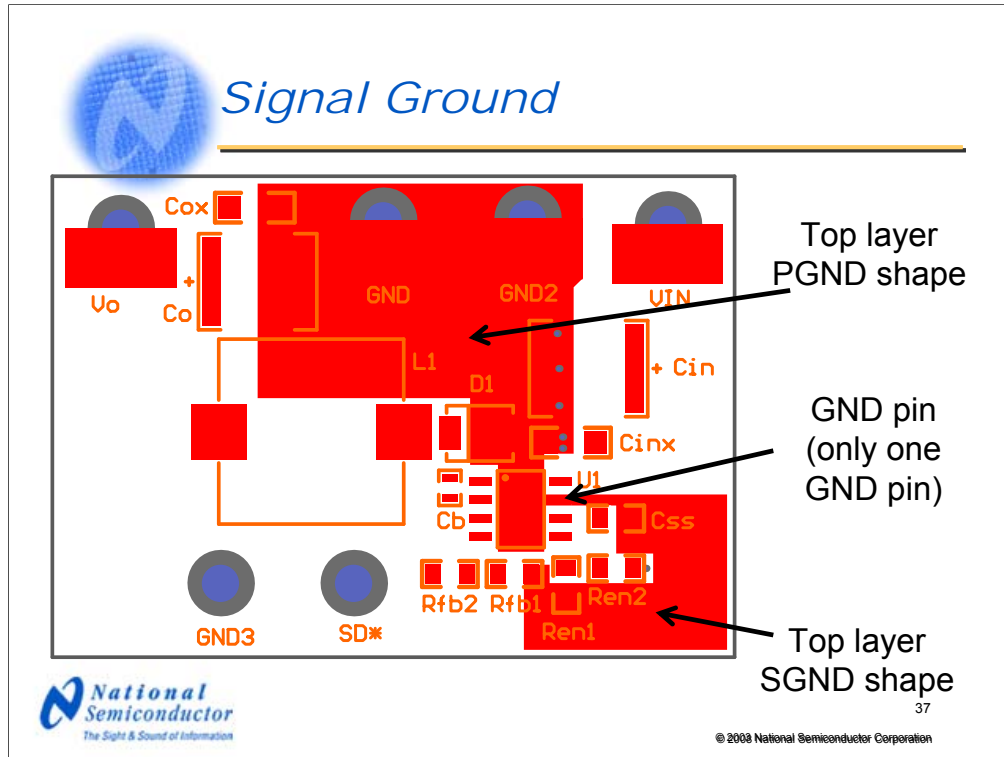
36

© 2008 National Semiconductor Corporation

Trace inductance in series with an inductor is not generally much of a concern. It will simply have the effect of increasing the total inductance in the path. In contrast, you don't want to add a significant amount of capacitance in parallel with an inductor. That's the analog of adding inductance in series with a capacitor and will cause problems. If you must add some inductance in series with an existing inductor, it's usually better to make the output path the longer one so you minimize the size of the noisy switch node.

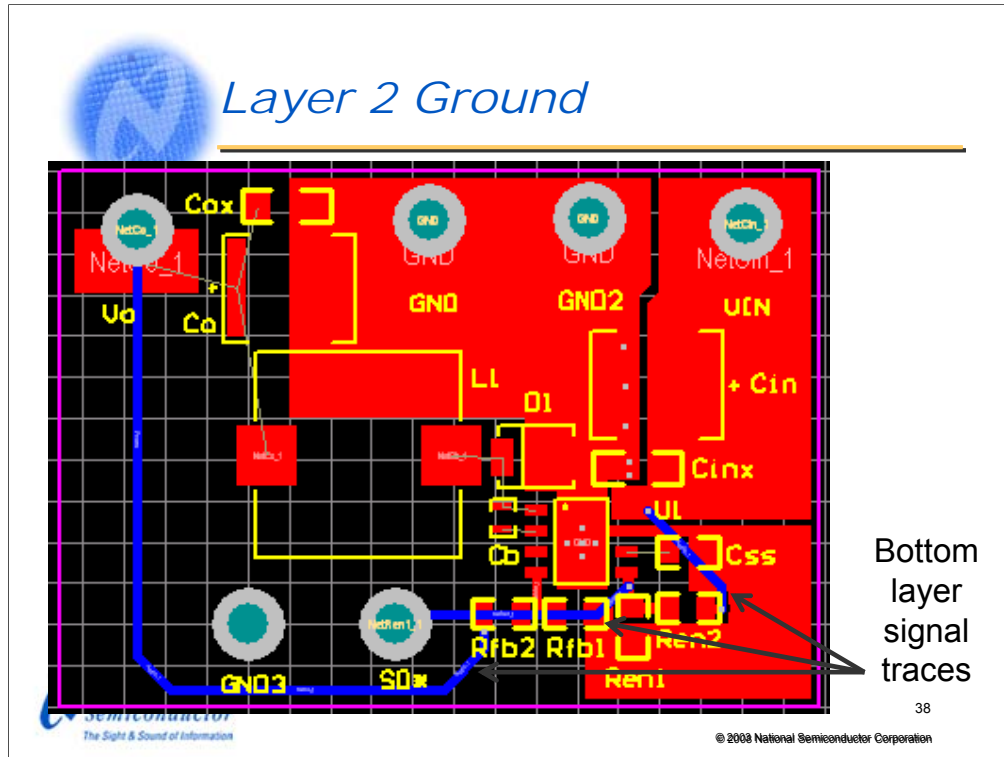
We've already discussed the use of vias to connect bypass caps to planes. A minimum of one via connected tangent to each pad of the capacitor is required. Two per pad is better, and three is slightly better but way down the diminishing returns curve. You can't have too many, but beyond two, they will do you very little added good.

The mantra through board layout needs to be "minimize stray inductance". Treat every high di/dt path as a high frequency RF connection (because it is!). You may be inclined to say "But I'm only running at 100 kHz", but the fast edges will have frequency components out to the tens of MHz to possibly hundreds of MHz. Don't ever forget that when doing a layout.



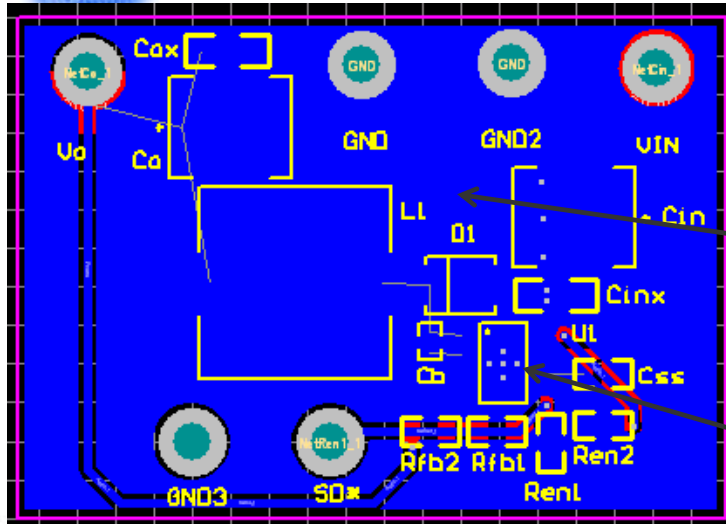
Since we've already decided that ground is a conductor with parasitics the same as any other conductor, it follows that if any current flows through this ground conductor, there will be an associated voltage drop. There's no place that you can point to and say: This represents the common reference point for all circuitry. That statement is true only for that one fixed location and, unless all ground connections return through zero current traces to this one point, there is no other location on your PCB that can be assured of being at the same potential.

The solution is to carefully partition grounds so as to know precisely where control signals get referenced relative to high-current ground paths. For control circuits like feedback dividers and soft-start capacitors, connect the ground referenced end of the part directly to the SGND pin of the control IC (if available). In this example there is only one pin, labeled "GND", so we connect the SGND net to that pin, in one place only. There may be large currents flowing through the main plane beneath and a significant associated voltage gradient. If you drop the bottom end of a feedback divider into a convenient point on the plane you may well introduce an error of many millivolts into your output voltage. It's more difficult to route a separate trace from the divider to the SGND pin, but there's almost no reasonable way to estimate the current paths through your ground planes, so assume they're a problem and design accordingly.



It's common to assume that ground is a solid, magical reference plane. But in many cases this tendency to ignore the realities of the ground scheme has led to tiny little connections left in the ground plane to carry very high currents or supposedly short return paths having to meander half way across a PCB to get back to their source. When there are lots of vias or traces laid in the ground plane you need to look very carefully and verify that it hasn't been nearly cut into pieces. Instead of a 2"-wide plane you may see a 20 mil-wide remnant after the plane gets cleared from around all the vias. On another note, take a look at the at the long trace from Vout to Rfb2 that allows us to minimize the trace to the FB pin of the control IC.

Layer 2 Ground



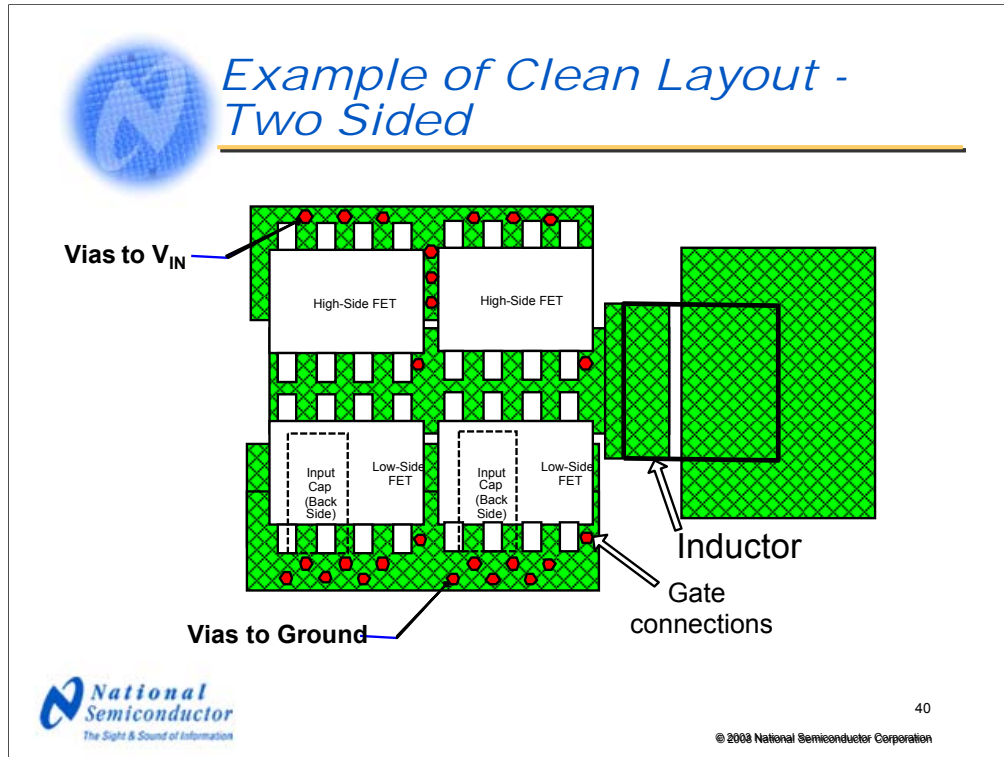
Large, unbroken area => low impedance, good heat transfer

Thermal vias from top GND to bottom GND

It's a very good idea to turn off all layers except ground at the end of a design and make sure you still have good plane integrity when you're done laying in all the signal and power traces and have dropped a collection of vias through the board. Vias are a necessary evil in many cases and should be avoided as current-carrying elements in the power path, if at all possible. The only time they become desirable is if they can be used to introduce redundant copper areas to a design, such as an inner layer being hooked in parallel to an outside trace area.

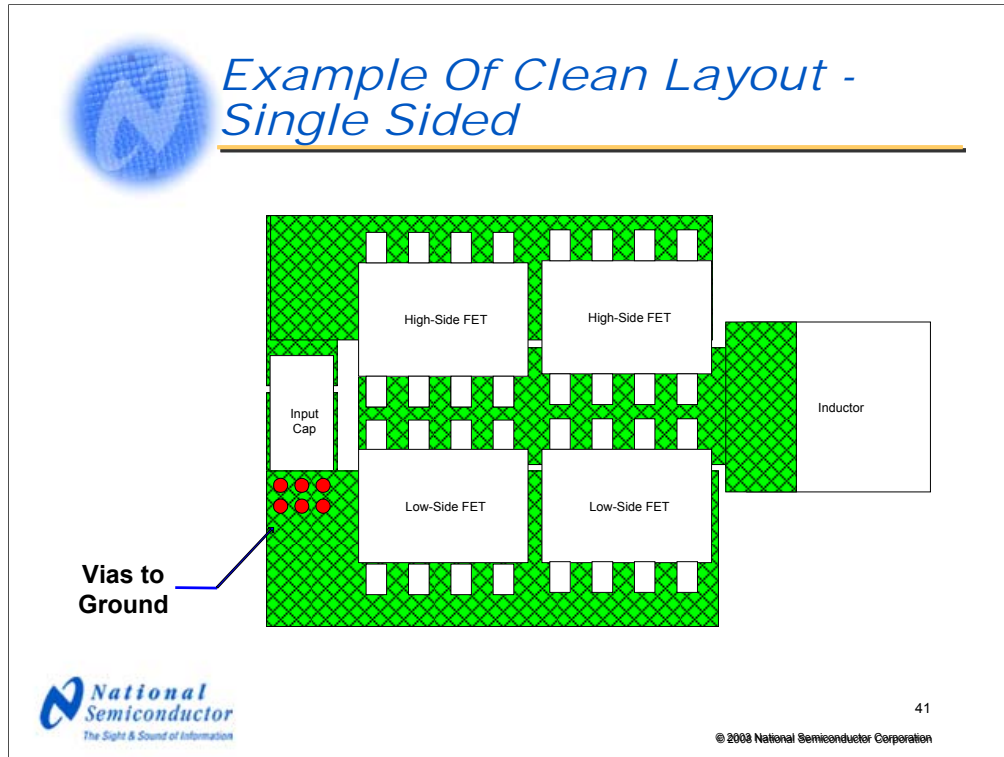
Vias are used as heat pipes because they have the ability to help conduct top-side generated heat to the back side of a PCB. The more vias that can be connected to a hot plane area, the more heat spreading can be achieved.

Something to watch for when using lots of thermal vias is that the internal ground planes can get badly cut up and end up with nearly no useful current path. This is particularly likely to occur on very small boards like voltage-regulator modules.

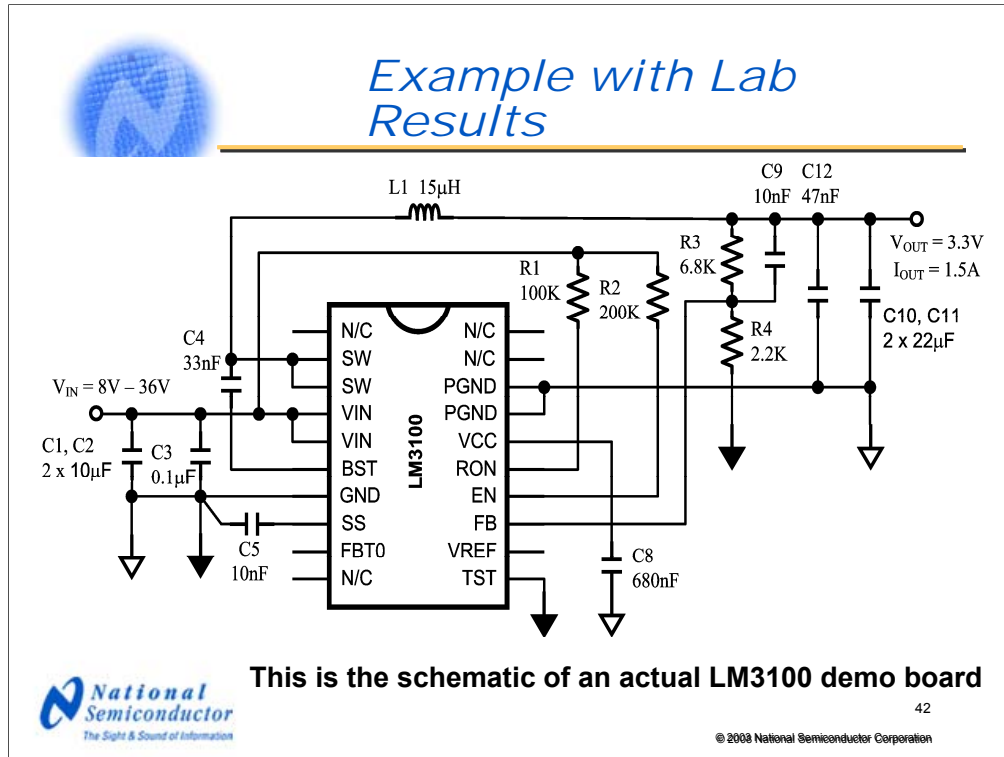


Here's an example of a reasonably clean layout for two parallel FETs, top and bottom, in a buck configuration. Note the locations of the decoupling caps on the back side. There would be a plane area between the cap V_{IN} terminals and the high-side FET drain connections. In practice, there needs to be small cutouts around the high-side FET gate pins. Install vias to inner layer gate traces along side the gate pins so as not to destroy the integrity of the switch-node power path. If you need to add more metal to the switch node, it will generally work out OK to double up the number of layers dedicated to this connection and add a number of vias close to the FET pins and near the inductor pad.

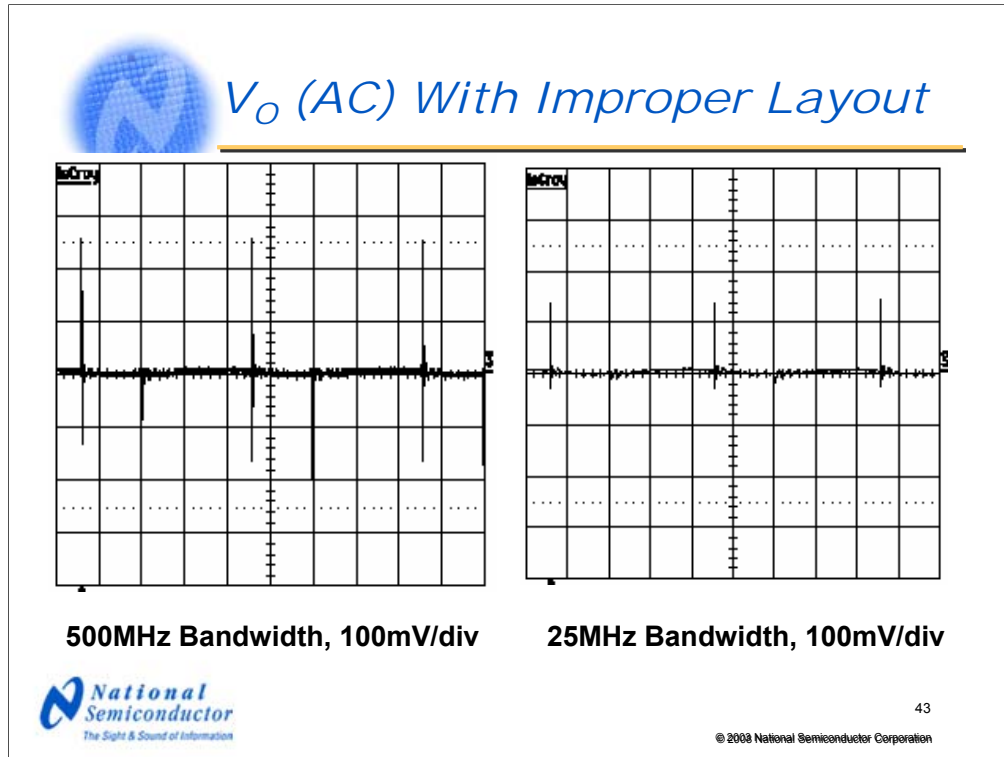
Note that the pad connected to the output side of the inductor is larger than the switch-node side. It's OK to run this island under the inductor and stop just short of the switch-node side pad. It's tied to the output and is an AC ground and, as such, will not inject noise. This is the output voltage, which should be nice and quiet, so it is one of the planes that can be extended for heatsinking purposes.



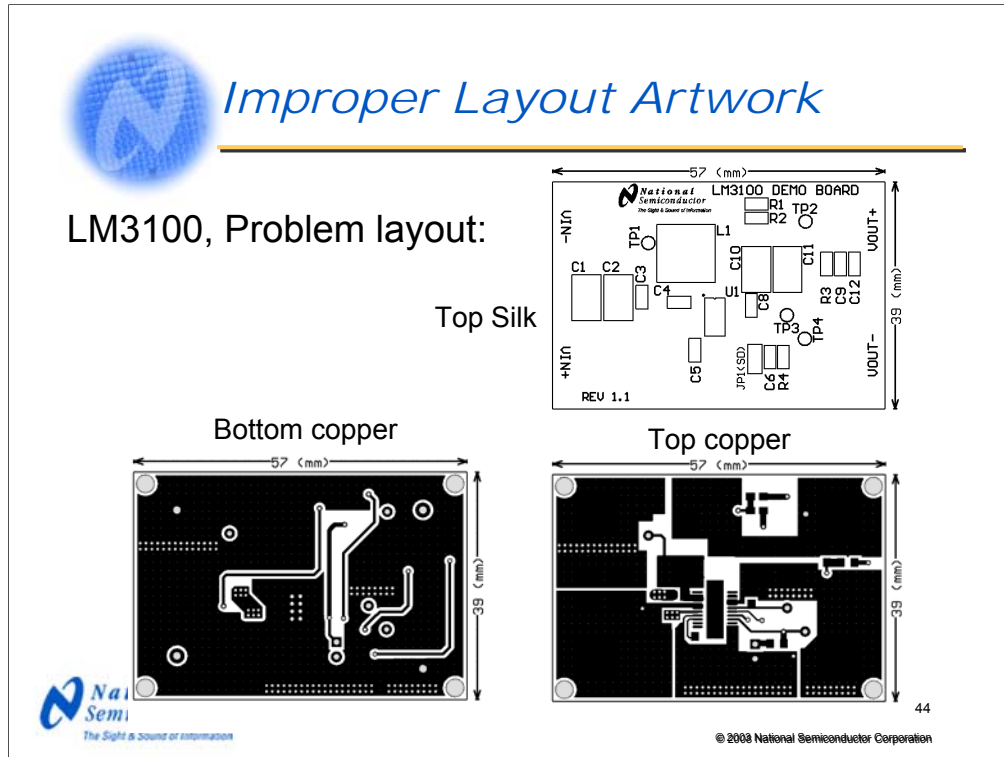
If you're forced to keep the input capacitor on the top side of the board, this approach is a decent alternative. Note that the ground vias are close to the cap ground, not at the FET-source connections. Otherwise the recommendations of the previous slide all apply here as well. Note that the input cap should be a high-frequency ceramic capacitor but has to be large enough to handle a significant amount of energy. Some thing on the order of $1\mu\text{F}$ or more for currents of several amps is appropriate at this location. Use $10\mu\text{F}$ or so for currents greater than 20A. The main-input supply-bulk capacitors can be located a short distance away. The local bypass needs only handle the very high speed edges of the FET currents.




The schematic above is the circuit used for the test results in the following comparison. No changes in component selection were made, only the layout of the PCB varies in the two examples. The design is a synchronous buck regulator with the two power switches and all the control circuits on the IC. Input for the tests that follow was 12V with a 3.3V at 1.5A output.



Shown here are two plots of output voltage AC coupled. The picture on the right was taken at 500MHz bandwidth while that on the left was taken at 25MHz bandwidth. Note the vertical scale for both shots is 100mV/div. So the spikes are quite large (nearly 300mV pk-pk) even over the limited bandwidth.



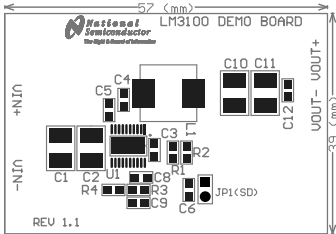
Here is the artwork for the problem design. Note that the PGND pins (18 and 19) connect to the OUTPUT cap ground rather than the input cap grounds. The input capacitor ground connections are located in the top left area of the board so the switch currents must return all the way around the board thru a large array of vias and then by way of the very cut-up ground plane on the back side of the board. That large loop is the primary cause of all those output spikes. A less significantly consideration is the relative location of the output terminals. They are located over 1" apart, which makes it difficult to make a good ripple measurement. Measuring directly across the output capacitors results in a slight improvement, but still no where near as good as the second layout option. Note also that the small signal components for the feedback divider etc. are spread out across the board. This is very undesirable.



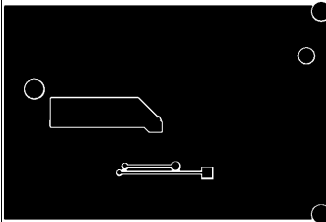
Improved Layout Artwork

LM3100, Proper layout:

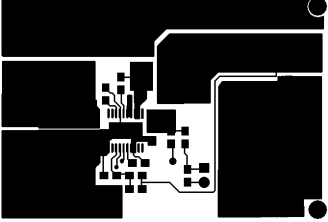
Top Silk




Bottom copper



Top copper

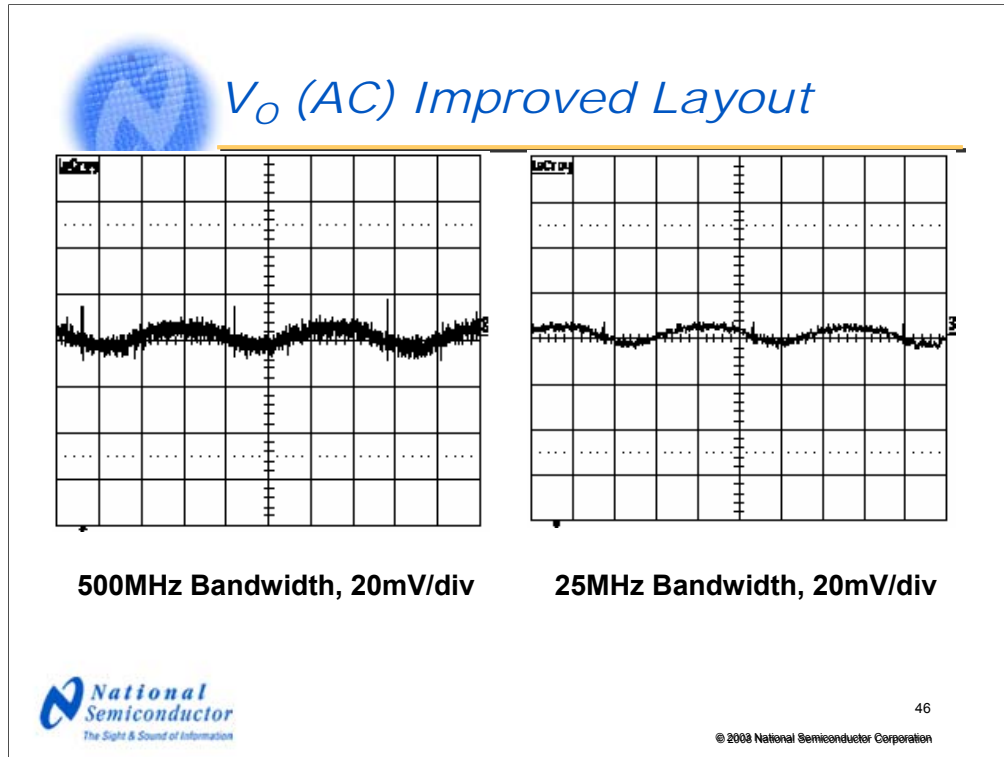




45

© 2008 National Semiconductor Corporation

Here you see the layout of the quiet board. Notice that the active circuitry actually takes up less overall space than the previous version. The input capacitors, C1 and C2 connect directly to the part's die attach paddle which is in turn connected to the PGND pins of the IC. This keeps the high di/dt currents out of the main ground plane. The small signal components are positioned close to the IC with short traces connecting to the appropriate pins. The two output terminals are located close together to permit good probing techniques to be used and minimize the loop area between the terminals and the output filter capacitor, C12. Since this is done on a two layer board there was the need to run a couple of traces thru the ground plane on the back side of the board. But notice that the runs are configured in such a way as to minimize the disruption of any return currents that have to flow on through the plane. Of course the top-side grounding was configured in a manner intended to virtually eliminate the need for AC currents to return through the plan in the first place. The result is a tight layout, that's extremely quiet.



All of the test conditions are the same for the results obtained above as for the improper layout. In fact, the bill of materials is identical! The only difference is in the PCB. Note also the dramatic scale change. These shots were taken at only 20mV/div. Even at 500MHz bandwidth the spikes are approximately 20mV pk-pk. With the bandwidth limited to 25MHz the spike amplitude is actually below the overall ripple envelope.

Again, the main difference in the boards is the way the low-side switch and the input caps connect.



Layout Considerations -Summary

- **Know where the high di/dt paths are in your design and minimize their loop area**
- **Use good grounding strategies**
- **Minimize Parasitic Inductance!**
- **Segregate signals and power**
- **Leave plenty of copper**
- **Add thermal vias where possible**
- **National Semiconductor application notes:**
 - **“Layout Guidelines for Switching Power Supplies,” AN1149**
 - **“SIMPLE SWITCHER® PCB Layout Guidelines,”**



47

© 2008 National Semiconductor Corporation

Here is a checklist of the layout considerations we discussed for noise reduction. Realize that PCB traces have impedance (both resistive and reactive) and current takes the least path of impedance.

