Optimizing PCB layout for HV CoolGaN™ power transistors

Practical guidelines to get the best operation of HV CoolGaN™ in switched-mode power supplies (SMPS) application

Authors: Eric Persson, Francesco Di Domenico

About this document

Scope and purpose

The fast-switching capability of high-voltage (HV) GaN transistors makes PCB layout challenging. This application note discusses several key concepts to help users understand the layout challenges and strategies to help optimize the layout for the best overall electrical and thermal performance.

Intended audience

Switched-mode power supplies (SMPS) design engineers, PCB layout engineers, technicians and developers of electronic systems interested in the optimal performance achievable with HV GaN devices.

Table of contents

1 Introduction

PCB layout has been an integral aspect of power electronic design since the first switching power supplies appeared more than 40 years ago. Regardless of the transistor technology, the parasitic impedances added to the circuit by the PCB layout must be understood and managed for the circuit to function correctly, reliably, and without causing undue electromagnetic interference (EMI).

Even though modern wide-bandgap power semiconductors do not suffer the severe reverse-recovery problems of older silicon technologies, their much faster switching transitions result in even more extreme commutation dv/dt and di/dt than their silicon predecessors. The common advice offered by generic application notes on power electronic PCB layout is to “minimize parasitic inductance as much as possible”. However, the best way to do that is not always clear. Moreover, not all conductive paths necessarily need to be the lowest possible inductance: consider the interconnection to an inductor – clearly there will already be inductance in that path.

It is of course impossible to minimize all interconnect inductance, and simultaneously eliminate all node-to-node capacitance on a PCB. The key to successful PCB layout is therefore to understand where the impedances really matter in switched-mode power electronics, and how to mitigate any undesired consequences of this inevitable impedance.

An additional complicating factor is that PCB layout not only involves optimizing the electrical interconnection, but often requires thermal pathways that conflict with electrical optimization goals. Even mechanical structures like heatsinks, when applied to the PCB, and separated only by a thin thermal interface material (TIM), can behave like an additional electrical plane of the PCB assembly, and interact with the switching nodes of the circuit.

This application note begins by explaining the fundamentals: what is really happening during a switching transition, what is the cause versus effect of the transient voltages and currents we see, and where exactly the current is flowing. When we think about current flow, we often forget to consider the return path, which is critically important. An additional key concept is how we think about inductance: it is often viewed as individual inductive elements that all add up around a loop. But they don’t necessarily all add up: depending on the geometric relationship between source and return current, the mutual inductance can change the sign, resulting in subtraction rather than addition. The concepts of loop, partial, and mutual inductance help us to explain and understand this interaction.

Next, different power-stage layout options are presented, along with the tradeoffs involved with each. The overall goal here is to understand the best ways to minimize power-loop inductance. With traditional through-hole transistors mounted perpendicular to the PCB, the inductance of the transistor package is independent of the PCB layout because they are at right angles. For SMT packages, the package inductance itself is a function of how the return path is routed, so there are more layout options and alternatives to improve overall performance.

Since design of the power loop includes thermal as well as electrical path optimization, the options and tradeoffs of top versus bottom-side cooled transistor packages are covered. Finally, the design, layout and routing of the gate-drive circuit, along with its “hidden” current paths, are explained.

2 The practical problem

The physical layout and packaging of power electronic circuits adds “parasitic” circuit elements R_Parasitic, C_Parasitic, and L_Parasitic. These parasitic elements can cause unexpected behavior and unintended consequences, circuit malfunction, EMI, oscillations and, in severe cases, cross-conduction or “shoot-through” that can lead to transistor failures. Resistive parasitics are comparatively easy to understand – especially for DC current. The solution to minimize parasitic resistance is to use more copper – to increase the total current-carrying cross-section. With high-frequency AC currents, the situation is more complex due to the skin effect. For PCB integrated magnetics, the skin and proximity effects need to be carefully considered, but that is outside the scope of this document.

The concept of parasitic capacitance is also straightforward. Especially in a structure like a PCB, where the copper layers form parallel plates with thin dielectric layers in between. We can use simple 2D tools to estimate the $C \approx \varepsilon_0 \varepsilon_R (area/spacing)$ and we can easily estimate capacitance per area for a given layer stackup. As we will see later, sometimes the capacitive coupling paths comprise components rather than the PCB itself. The question which capacitance value is acceptable will be covered later as well.

Parasitic inductance is different: basic circuits classes teach us to think of inductors as discrete elements that sum like resistors in series. However, in more advanced magnetics courses, we learn they interact with each other through mutual inductance, which can either increase or decrease the total inductance, depending on the geometry and direction of current flow. Also, we often don’t have a good estimate of layout inductance, or know the magnitude of dI/dt to expect from our switching circuit – how much will cause problems?

These layout issues are not new to power electronics, but GaN transistors with low charge and no reverse recovery make switching transitions even shorter. Fast-switching transistors primarily cause two interrelated issues. The high transconductance of GaN, combined with its low gate charge, can result in extremely fast switching $(dI_{DS})/dt$. The fast $(dI_{DS})/dt$ leads to high peak currents as the capacitance of the switch-node is rapidly discharged – the resulting C $(dV_{DS})/dt$ adds to the load current. And high peak switching current combined with the low Q_{OSS} of GaN power transistors causes those fast $(dV_{DS})/dt$ edges at turn-on.

Why are fast d*I*/dt and d*V*/dt problematic? On the one hand, fast turn-on d*I*/dt reduces switching time and therefore loss, so it should be desirable. But the problems that occur are primarily due to the L d*I*/dt reaction voltage that appears across parasitic inductive elements. Most commonly the undesirable effects occur in the main commutation loop, or in the gate-drive loop. Voltage overshoot in the power loop can increase EMI issues and create HV stress on the transistors, leading to reduced reliability. In the gate loop, L d*I*/dt reaction voltage will slow switching speed by subtracting from the applied gate voltage, but can also lead to ringing and overshoot of V_{GS}, or even oscillatory behavior, which can quickly destroy the transistor.

3 A detailed look at the half-bridge in hard-switching

The half-bridge topology is widely used in power electronics. It is the basis of “totem-pole” bridgeless power-factor correction (PFC), full-bridge DC-DC converters, LLC converters, inverters, and many others. Because GaN has low output charge and no body-diode recovery, it is an ideal transistor candidate to use in a half-bridge, and it can be used interchangeably in either hard- or soft-switching. Consider the half-bridge as a two-port network, as shown in **Figure 1**.

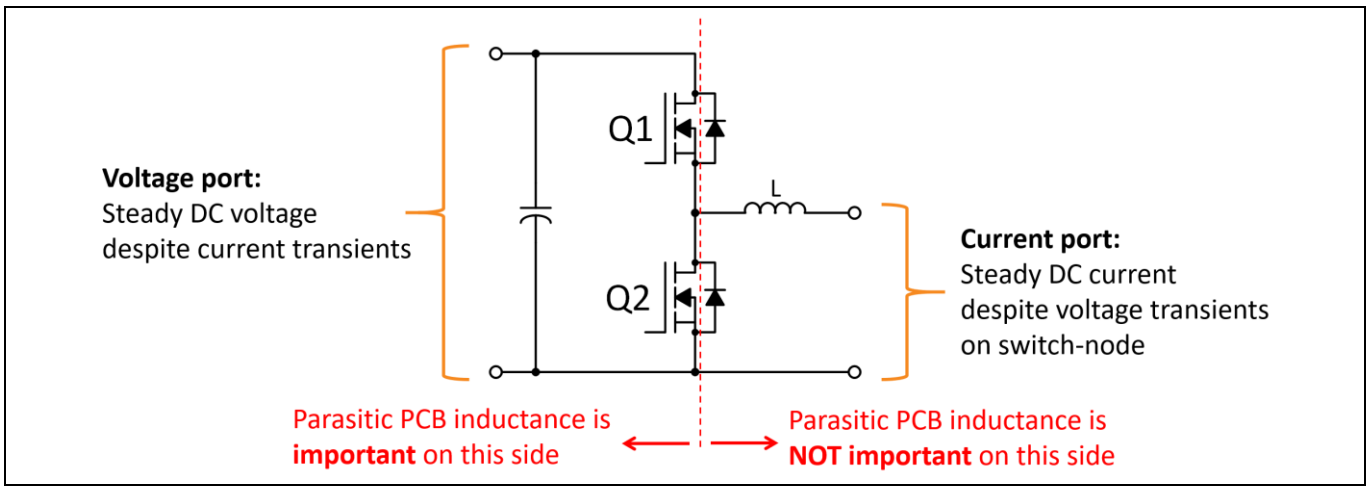

Figure 1 The half-bridge topology is a two-port network

The left side of the circuit is a voltage port that represents a DC bus. The DC bus voltage should remain steady despite the current transients that occur when the transistors switch. It doesn’t matter if the power flow is to the left or right; in either case, the bus should remain fixed. The right-hand side of the circuit is a current port. Here, the inductor current should remain steady despite the voltage transients on the switch-node. This is an important distinction, because it indicates that parasitic inductance is important for the loop on the left side – to minimize $L\,dI/dt$ transient voltages from appearing across the transistors. But the right side is already an inductive path – so here it really doesn’t matter if some extra parasitic inductance is added – it will be insignificant compared to the intended inductor.

Another useful concept to help understand switching behavior is to think about what the forcing function is. Rapidly changing currents cause a reaction voltage $v = L\,dI/dt$, and rapidly changing voltage induces current $i = C\,dV/dt$, but which is the cause versus effect? For hard-switching transistors, current is the forcing function. During switching, the transistor acts as a transconductance amplifier – driving current in response to the gate signal. How quickly the transistor turns on (dI/dt) is limited by how quickly the gate can be fully enhanced. dV/dt, on the other hand, is an effect: it is the result of how rapidly the applied current can charge the node capacitances.

Now consider the switching transients in the half-bridge. **Figure 2** shows a half-bridge setup for pulse testing: Q2 is the active switch, and Q1 serves as a synchronous rectifier. The switching diagram shows the hard-switched turn-on of Q2, followed by its turn-off (which is essentially zero-voltage switching, ZVS). This emulates the typical circuit operation of a totem-pole PFC operating in continuous conduction mode (CCM). In both cases, the initial inductor current is 10 A, and the bus voltage is 400 V.

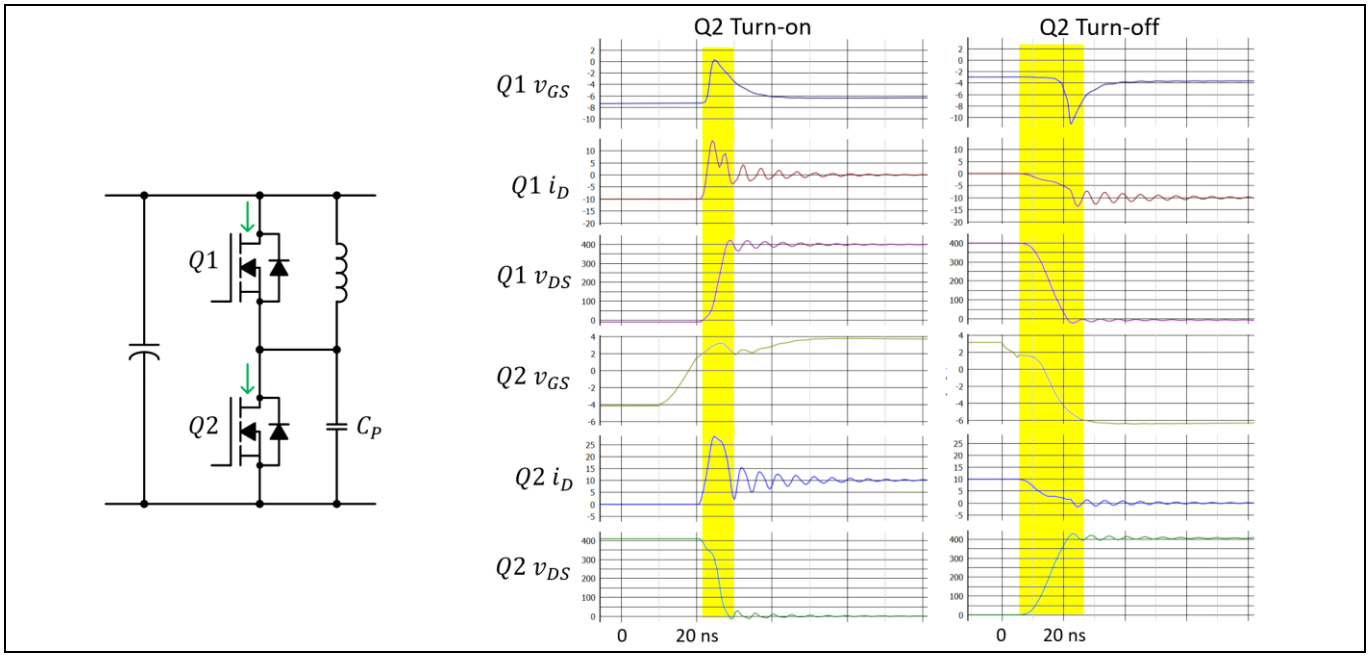

Figure 2 Hard-switched turn-on and turn-off of a GaN half-bridge

Note that when Q2 turns on, its drain current far surpasses the 10 A inductor current, peaks at 28 A, then returns and settles down to the 10 A inductor current after the resonances damp-out. This waveform looks suspiciously like reverse recovery, but the key is to look at the timing of v_{DS} versus i_D. Reverse recovery prevents the high-side diode (Q1) from blocking voltage – so it essentially remains on (conducting), until the peak of the i_D waveform – *then* v_{DS} begins to change. That is the classic signature of reverse recovery. But here, v_{DS} is clearly beginning to move as soon as the drain current exceeds the inductor current, indicating that the response is purely capacitive with no reverse recovery.

The area under the 28 A peak, above the 10 A inductor current line, represents charge that Q2 has to discharge. It is the sum of Q1 Q_{OSS} plus the parasitic PCB capacitance represented by C_P. These currents are very difficult to accurately measure on the PCB unless a dedicated wide-bandwidth current-shunt is added to the circuit. Here we use simulation to estimate the current, and a value for C_P is added to the simulation circuit. But even with this simulation, what we don't see is the internal self-discharge current of Q2 Q_{OSS}. While it does contribute to hard-switched losses, the discharge path is entirely contained on the transistor die, so the effect of this added capacitive current is not seen in the PCB layout.

Note: The transient current during switching is difficult to measure, and simulation requires adding accurate parasitic elements. In this case, the peak current is 18 A above inductor current, and the slew rate is about 9 A/ns at turn-on. This represents a typical value measured on dedicated GaN test setups and characterization platforms, where the turn-on dI/dt is typically measured in the range from 4 to 16 A/ns. This typical value will be used in various layouts to assess what the resulting transient voltages could be.

4 Mutual and partial inductance

Before addressing how to minimize parasitic inductance on a PCB layout, a few inductance concepts will be reviewed. Consider a piece of wire formed into a single-layer loop 100 mm in diameter. Measuring the inductance of that loop with an impedance analyzer would indicate about 250 nH. You could make a piecewise linear model of this by connecting n small inductors in series around a loop, as shown in **Figure 3a**.

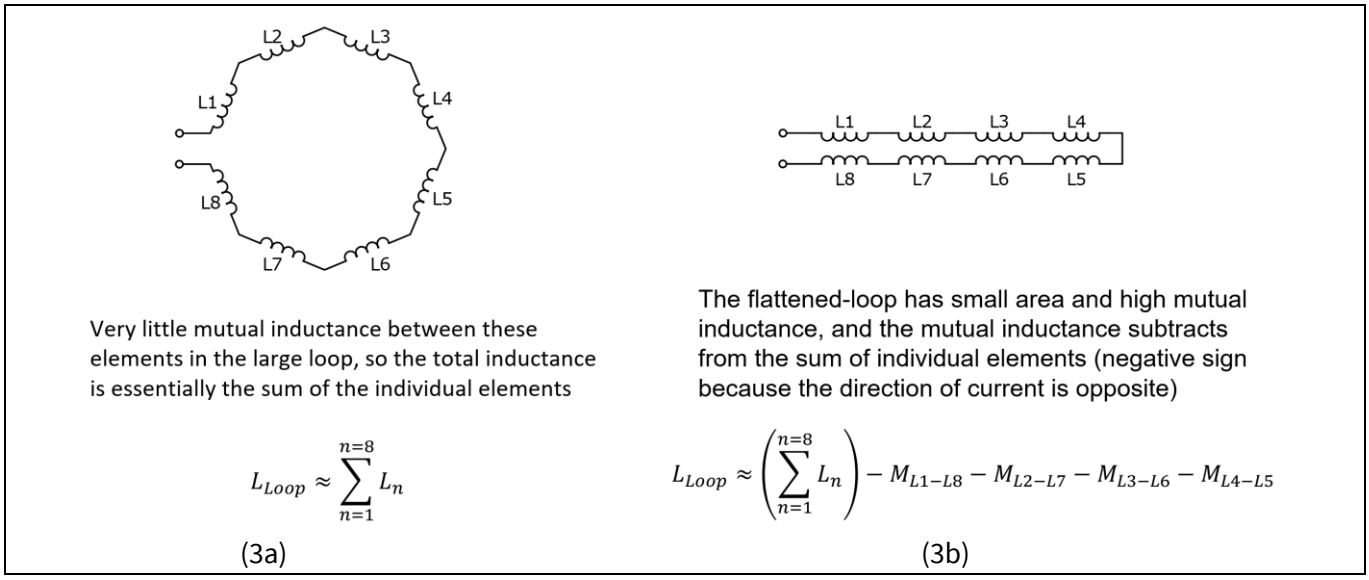

Figure 3 a) Model of inductances around a circular loop; b) the same loop, flattened

In this example, n = 8 for simplicity. Knowing that inductors in series sum just like resistors, each of the eight elements represents a 31 nH segment of the circle. This doesn't hold true if we take the same piece of wire and flatten out the loop as shown in **Figure 3b**: now the measured inductance drops by a factor of four or more (depending on the insulation thickness).

The change is explained by the concept of partial inductance **[1]**. The value of each segment L_n is equal to its *intrinsic* inductance in free space (its partial inductance), plus the effect of mutual inductance *M* from nearby segments. Note that mutual inductance has both a magnitude and a sign, so it can add to or subtract from the intrinsic inductance. In the example of **Figure 3b**, the direction of current flow in L1 and L8 is opposite, thus the mutual inductance subtracts from the partial inductance of each segment. The same holds true for all pairings shown. Power electronic engineers commonly think about mutual inductance when specifically designing a transformer, for example, but it seems much less obvious when simply thinking about conductive tracks on a PCB.

Another equally valid way to consider round versus flat-loop inductance change as shown in **Figure 3** is to apply Ampere's law and note that the loop *area* has been reduced, and thus the flux linked is also reduced, leading to the inductance reduction. These are simply two different ways to think about the same situation, and both are valid.

Consider the same wire loop experiment in the context of a PCB layout as shown in **Figure 4**. The loop begins at Cbus+ and ends at Cbus- (the +BUS and GND power loop). The individual inductive elements represent the wire bonds in the transistor package, the leadframe, copper segments on the PCB and so on. Comparing the two different layer spacings shown in **Figure 4** and **Figure 5**, it should be clear that the mutual inductance will be greatest when the spacing between the surface current and the return path current is smallest. This is also consistent with the smallest loop area.

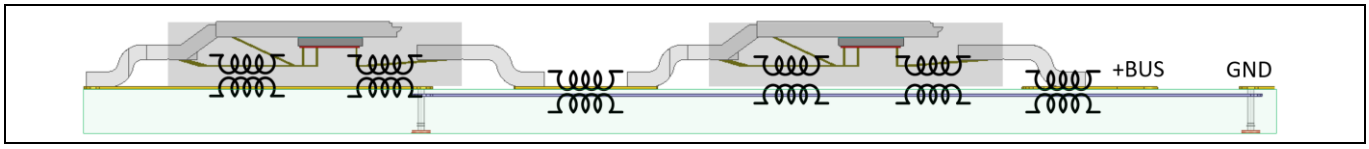

Figure 4 **Power loop with narrow spacing between layers maximizes mutual inductance**

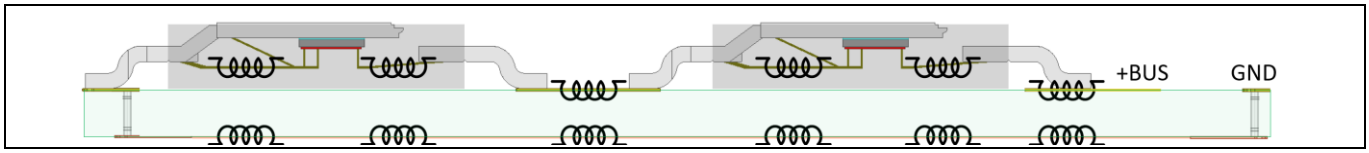

Figure 5 **Increased layer spacing reduces mutual inductance, resulting in a higher overall total power-loop inductance**

5 Package inductance: fixed value or layout dependent?

Generally, the calculated transistor package inductance is simply the partial inductance – no mutual inductance is presumed external to the package. In other words, the specified package inductance assumes the return path is infinitely far away and does not affect the partial inductance. For through-hole packages like the TO-220 and TO-247, this assumption is valid, since the package is commonly mounted perpendicular to the PCB and the current paths are orthogonal (thus no mutual inductance). But for surface-mounted packages, an optimized current return path will *significantly* reduce the effective inductance of the in-situ package.

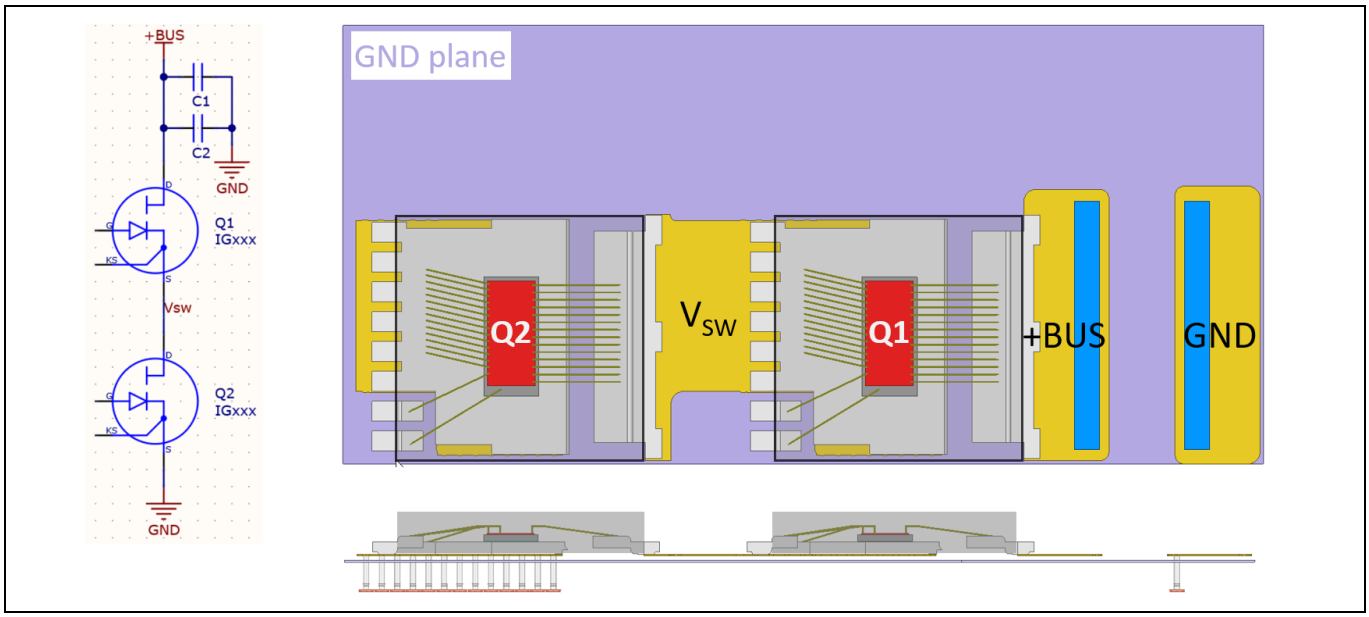

Figure 6 Low power-loop inductance using surface-mounted "TOLL"-packaged GaN transistors – L = 2.8 nH

To illustrate this, consider a simple half-bridge layout as shown in **Figure 6**. The power loop runs from the +BUS pad, through Q1 and Q2, from right to left through the transistor packages and on the surface layer (gold color). Then the current drops to the second layer (violet color) through vias below the Q2 source. The return current then flows in the opposite direction back to the GND on the far right. This is like the flattened loop in **Figure 3b**, where each of the parasitic inductances (a bondwire, for example) has a corresponding segment in the ground return plane.

The mutual inductance (or small loop area) makes the total loop inductance quite small. In fact, using 3D Finite Element Analysis, the calculated loop inductance of the geometry shown is only 2.8 nH. However, the total package inductance in the transistor models shows 2.1 nH per transistor. This is a perfect example of how the package partial inductance can be misleading, because it ignores the effect of the return path as discussed earlier. The explanation is that the 2.1 nH package inductance is reduced by the mutual inductance of the return path when laid out as shown here. While the loop inductance of this example is impressively low, there is a fundamental problem in trying to use this layout in a real power converter. There is no space to insert a thermal via field to remove heat from Q1 without creating a big hole in the ground return path under Q1. **Figure 7** illustrates the problem.

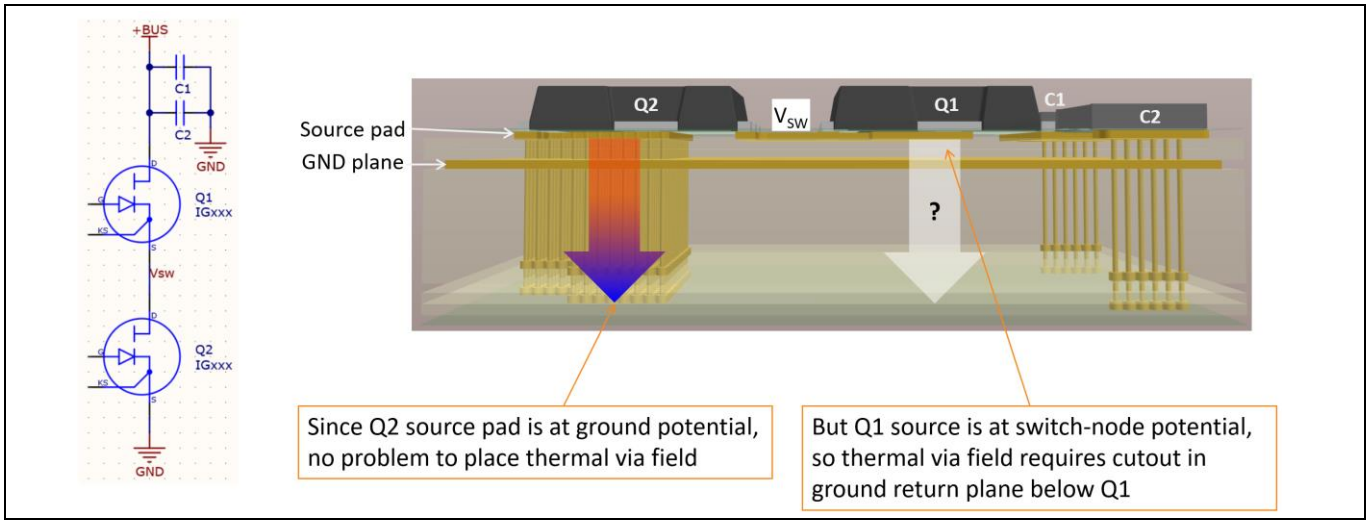

Figure 7 The need for the Q1 thermal path will make optimal electrical layout challenging

Adding a cutout for thermal vias will redirect the return current under Q1, and this will not only reduce the mutual inductance, but it will create a lateral loop with additional inductance. **Figure 8** shows the addition of the thermal via field, and the cutout necessary for clearance around the voltage of the switch-node.

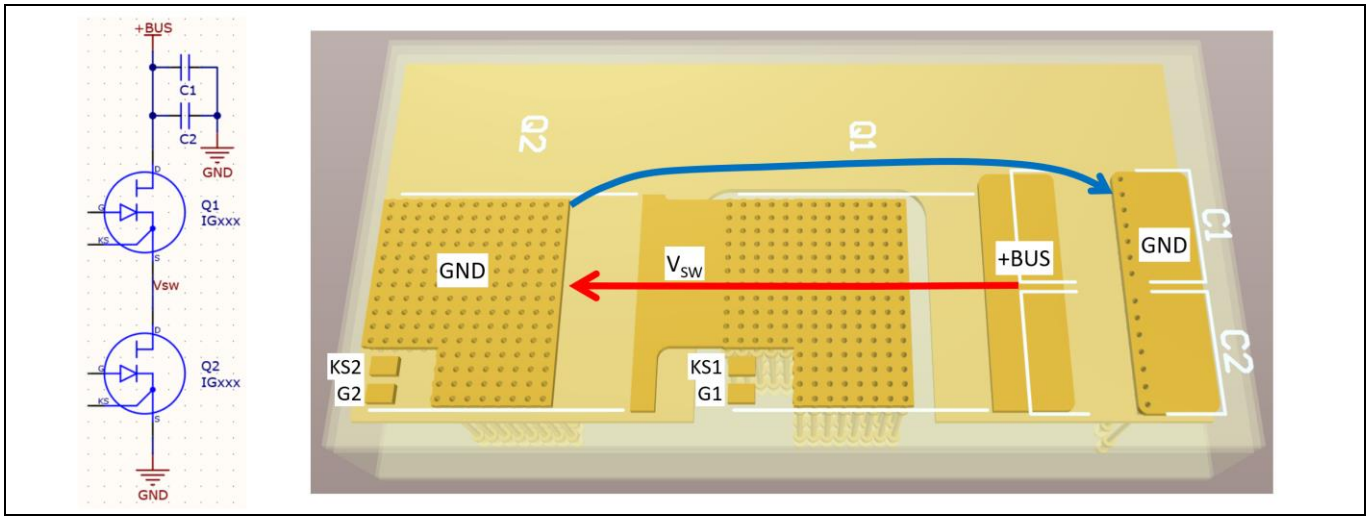

Figure 8 The ground-plane clearance around the via field creates a lateral loop in the return path – L = 8.8 nH

The current path through the transistors (in red) is essentially the same as before. But now the ground return path (in blue) is not directly below the red path, and thus a lateral loop is formed, and this reduces mutual inductance, increasing the total loop inductance. Comparing the two otherwise identical layouts: in **Figure 6** with no cutout, the loop inductance was 2.8 nH; in **Figure 8** with the cutout, the loop inductance more than triples to 8.8 nH.

The single-sided return in **Figure 8** does leave the gate pads open and easy to connect on either side of the PCB. What if a second, parallel ground return path was added on the gate side of the transistors? This example is shown in **Figure 9**. Now there are two lateral loops in parallel, so the parallel combination should reduce the lateral-loop inductance – and it does. The double-sided return path of **Figure 9** has a total loop inductance of 6.2 nH. This is still more than double compared to **Figure 6**, but a significant improvement over the single-sided return path.

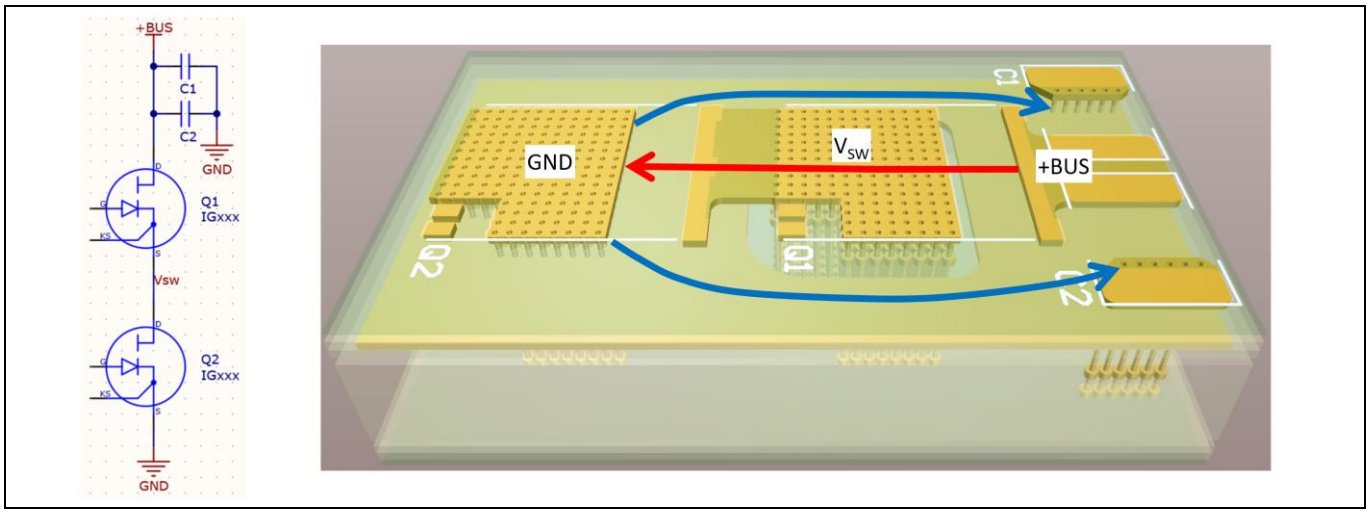

Figure 9 **A double-sided return path around the thermal via field helps minimize lateral loop inductance – L= 6.2 nH**

6 The advantage of top-side cooled transistor packages

The previous layout options using bottom-side cooled packages all require compromises in electrical layout to accommodate the thermal vias necessary to remove heat from the transistor(s). Adding 100 or more thermal vias per transistor not only compromises the electrical layout, but vias cost money: every drill operation (especially for small-diameter drills) adds cost to the PCB manufacturing process.

Another option to consider is using top-side cooled transistor packages. Sometimes these are simply the same bottom-side cooled devices with a flipped lead bend. But in most cases, the top-side cooled transistors are packages specifically designed to optimize both thermal and electrical performance of GaN transistors. **Figure 10** shows an example of the same half-bridge layout, but this time using a TOLT package.

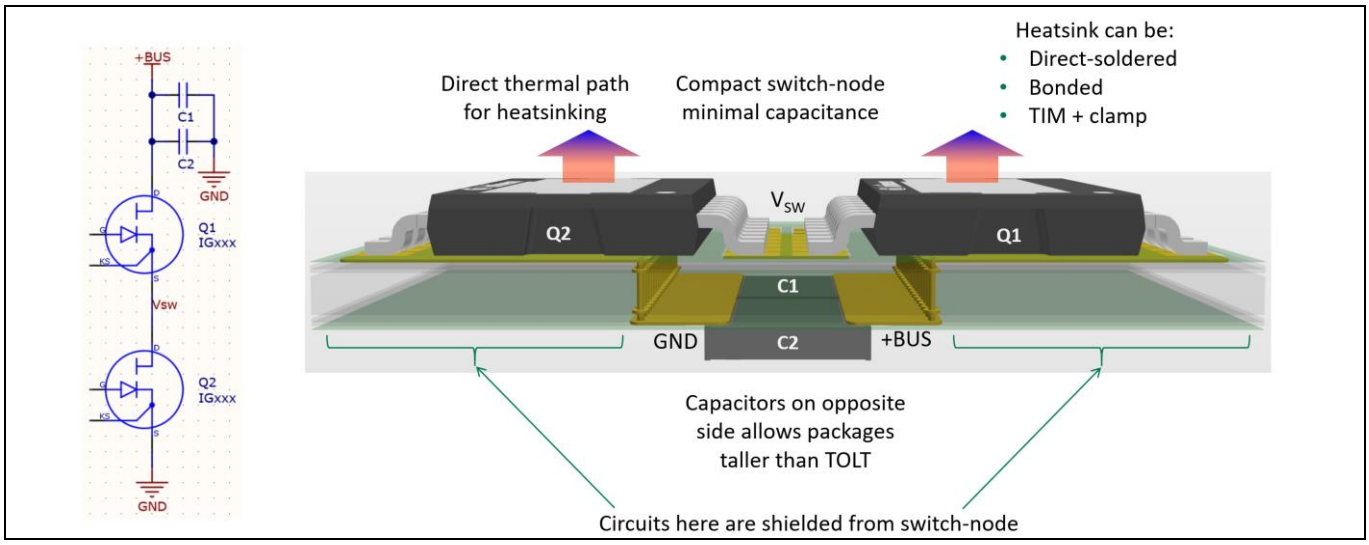

Figure 10 Top-side cooled transistors enable optimization of both electrical and thermal paths, resulting in 5.8 nH loop inductance

The big difference here is that, compared to the previous layouts, there are NO thermal via fields required below the transistors. This saves cost, and allows electrical layout optimization *independent* of the thermal path. An additional benefit is that the ground and +BUS planes, which are electrically "quiet" equipotential planes, serve as Faraday shields between the noisy switch-node and any other circuits on the bottom-side of the board. In this example, the high-frequency bus capacitors C1 and C2 are located on the bottom-side of the board. This is sometimes necessary, depending on the height of the capacitors, and the geometry of the heatsink, to avoid interference and enable proper creepage and clearance for the heatsink on top of the transistors. This option does leave a loop above the capacitors that increases loop inductance, but using a thinner board (0.8 versus 1.6 mm thick, for example) will help reduce that added loop inductance.

As another option, if there is room to locate C1 and C2 on the same side of the PCB as the transistors, then the layout in **Figure 11** can provide an excellent low-inductance power loop. Even though the return path is extended longer than **Figure 10** laterally, the spacing between layers is much thinner (0.18 mm, for example), making the overall loop slightly better than **Figure 10**.

The example top-side cooled packages shown here do not have an electrically isolated heatslug – it is a part of the source leadframe, so it is electrically connected to the source potential. The metal is solderable, so individual copper heatsinks could be directly soldered to each transistor. Alternatively, thermal interface materials can be used with many other heatsink options as well.

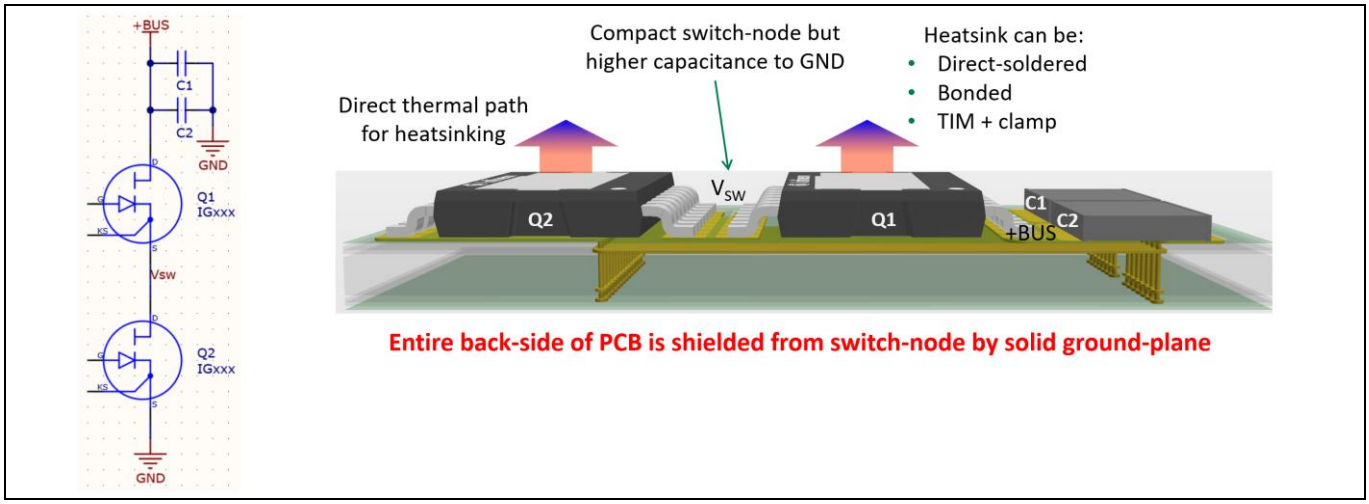

Figure 11 **Moving the capacitors to the same side as the transistors for lowest overall inductance provides 4.9 nH loop inductance**

7 Summary of power-loop layout options and results

All of these SMT layout options can be compared to a standard TO-247 package to see the overall loop inductance and estimate what the overshoot voltage would be for a half-bridge assuming 9 A/ns dI/dt, as previously discussed. The best-case TO-247 layout results in about 15 nH of total loop inductance, as shown in **Figure 12**.

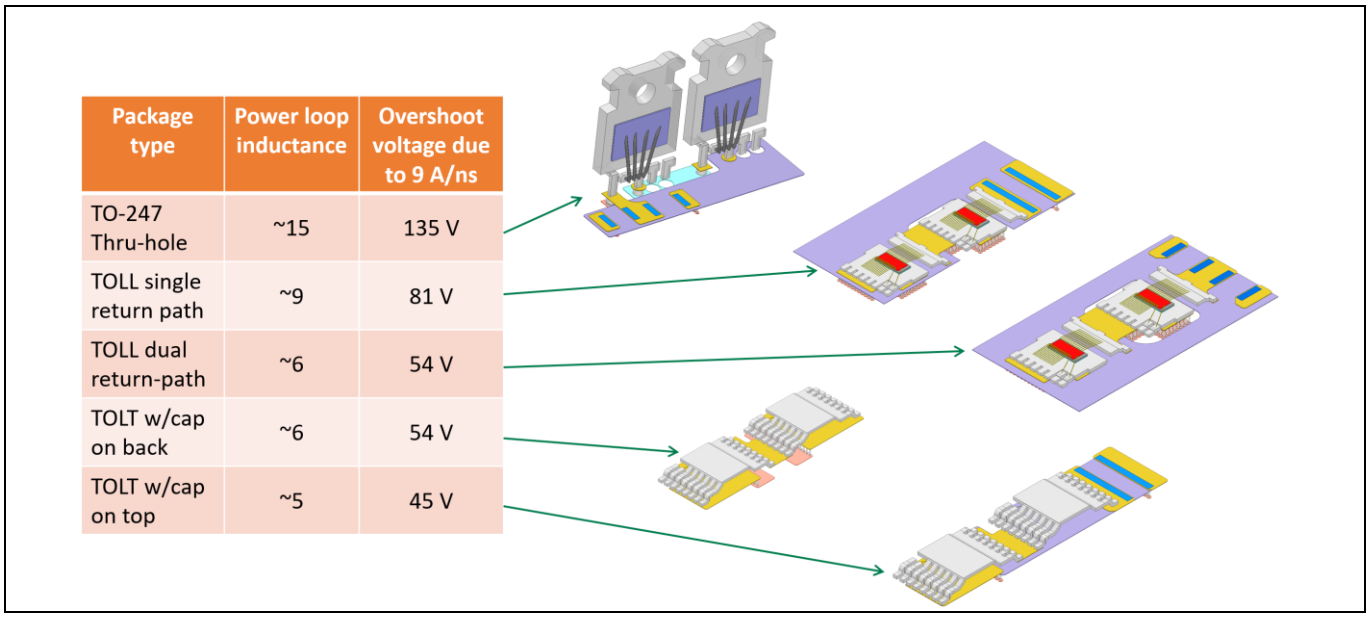

Figure 12 Summary of power-loop inductance and implications

Applying this knowledge to the expected di/dt, we see that the resulting overshoot voltage would peak at 135 V. This may exceed the design limits: for example, if the nominal bus voltage is 400 V, and the design rules mandate that peak voltages are less than or equal to 80 percent of rated voltage, then the 535 V peak will exceed that even for a 650 V rated transistor. This suggests that the solution for using TO-247 packages and keeping overshoot voltage below 480 V is to slow down the switching, by increasing the turn-on and turn-off gate-drive impedances, for example. Slowing down switching will of course also increase switching loss – which takes away from the benefit of using GaN transistors in the first place.

The surface-mount TOLL with single-sided return path is the next best option for low power-loop inductance. **Figure 12** shows that its layout inductance with 9 A/ns applied would result in an 81 V overshoot – just at the design goal of 480 V peak assuming a 400 V bus. By adding the parallel return path on the gate side (the TOLL dual return path), the loop inductance is now low enough that the overshoot voltage is 54 V, providing some additional margin so that even if the bus is pumped up to 420 V, the added 54 V overshoot will stay below the 480 V design goal. All three of the bottom layouts in **Figure 12** have sufficiently low loop inductance, providing suitable options for either top- or bottom-side cooled packages.

8 Considerations for gate-drive layout

GaN transistors have low threshold voltages, typically in the range of 1 to 2 V. In addition, the fully-on V_{GS} is in the range of 3.5 to 5 V (depending on the gate technology), and the transconductance of the transistor, as well as its gain-bandwidth, are quite high in the active region. This set of characteristics makes it imperative that the gate-drive loop must be low-impedance, otherwise $C_{GD}dV_{DS}/dt$ current injected through the "Miller" capacitance (C_{GD}) will influence the gate voltage, resulting in ringing, overshoot, potentially high-frequency oscillation and spurious turn-on, leading to potentially destructive cross-conduction or "shoot-through".

One of the biggest challenges is keeping the gate off when a fast-rising dV/dt appears on its drain voltage. It is all but impossible to make the gate-drive loop low-enough impedance with separately packaged transistor and driver. This is primarily why negative gate-bias is used in discrete designs: to provide sufficient margin so that gate bounce voltage does not exceed the threshold during switching transients.

We can use the concepts discussed earlier regarding power-loop inductance to optimize the gate-loop as well. **Figure 13** shows an example layout of a gate-drive IC connected to a TOLL-packaged transistor. Just like in the optimized power loop, the concept here is to use a pair of PCB layers with close spacing to maximize the mutual inductance (minimize loop area and gate-drive loop inductance). This example includes the RC network used with the GIT version of the GaN HEMT, including separate R_{ON} and R_{OFF} for the separate source and sink pins on the driver.

The six-pin gate-driver package U1 is on the far right with its supply bypass capacitor C3. The four RC components are in line with the gate pin of the GIT, all on the surface layer (red). The KS pin of the GIT is the reference point for the return path, which is defined by the copper-pour polygon on layer 2 (dark brown color). The return path terminates back at the "GND" connection of U1 and C3. **Figure 14** shows the routing more clearly without the component bodies including the vias that connect the KS pin as well as GND for the driver to the plane. This method of placement and routing for gate drive provides the best overall performance for designs using discrete gate-drive components (not integrated into the transistor package). Note that the gate-drive return plane uses layer 2, the same as the power-loop return plane to the left. Thus, implementation of this gate-drive layout can be used only with the single-sided return path of the power loop described earlier (because the gate-drive return plane occupies the same space needed for the additional power-loop return path). If a lower-inductance power-loop layout using the double-sided return is needed, an alternate layout of the gate driver is necessary.

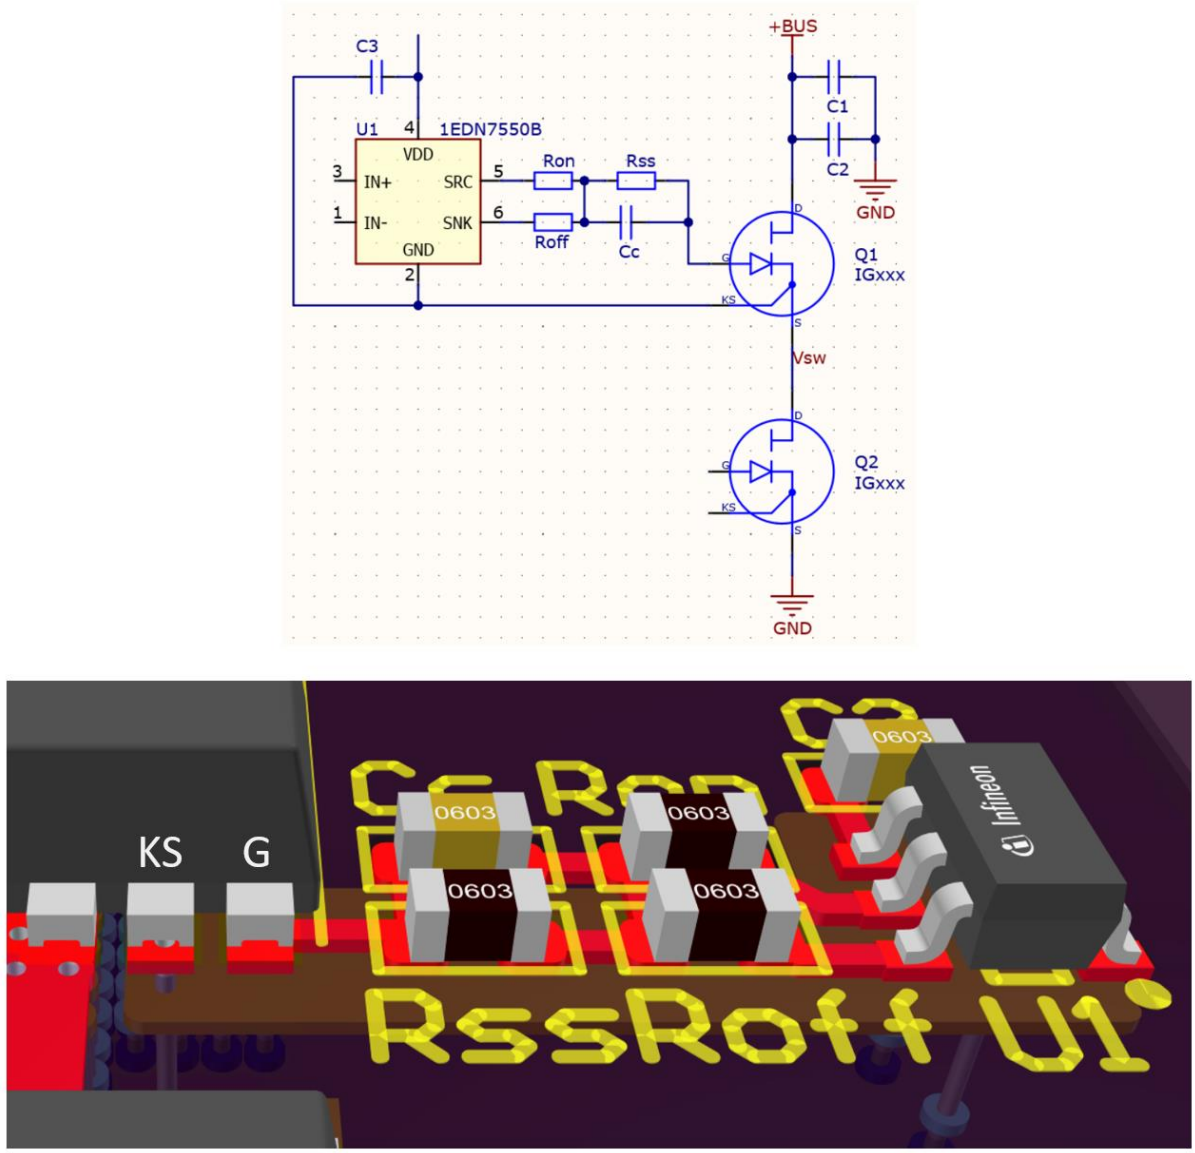

Figure 13 Example gate-drive layout with return plane directly below the drive circuit

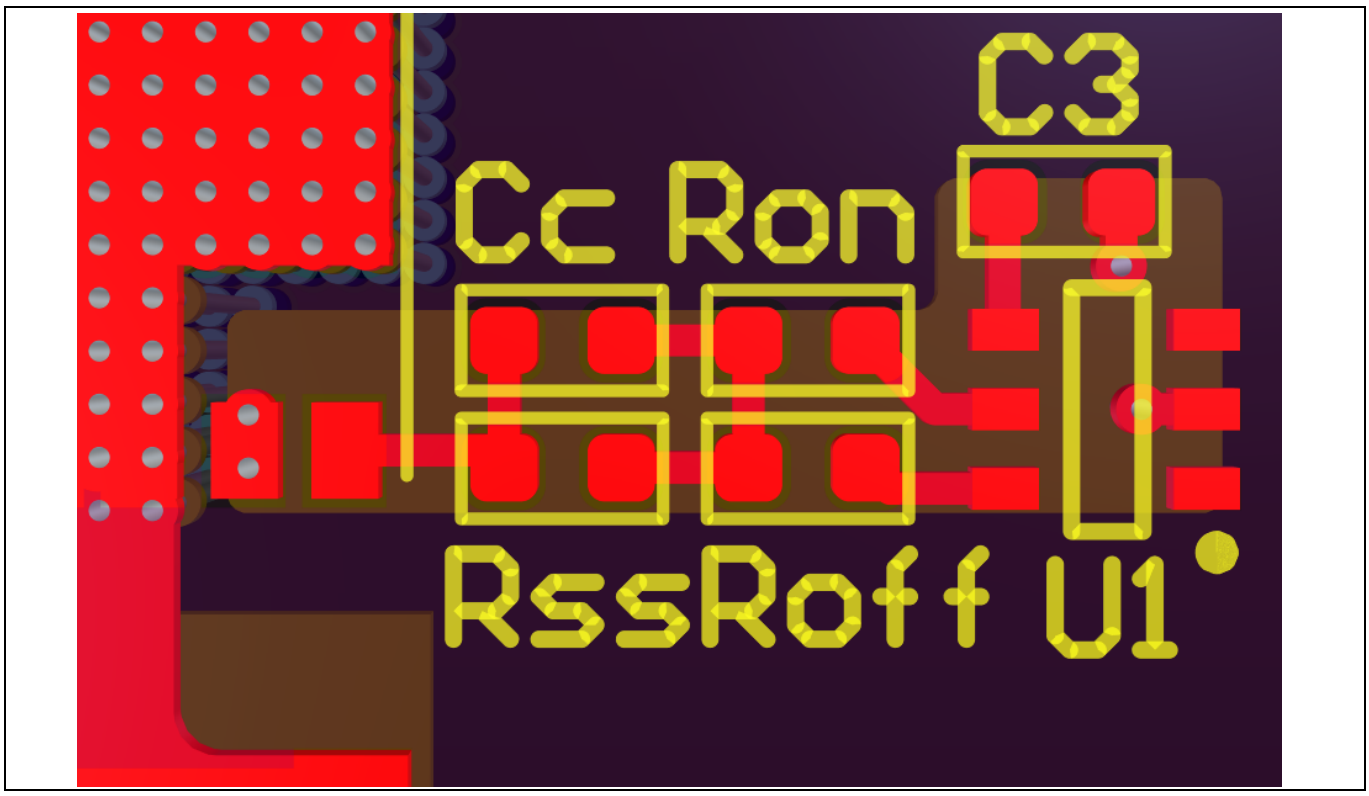

Figure 14 A more detailed look at the gate-drive loop routing from Figure 13

To keep the gate-drive circuit on the top layer next to Q1 and use the double-sided power-loop return path from **Figure 9**, there are two problems: first, since layer 2 for the return plane can't be used, a lateral gate-drive loop – all on the top layer – must be used as shown in **Figure 15**. Note that in this example, the gate-drive circuit has only turn-on and turn-off resistance for the Schottky gate HEMT, instead of the RC network of the previous GIT example. To keep the loop as small as possible, the return path is routed right next to the drive path (note how the KS pin connects to the driver pin 2).

The second problem is much more challenging: the high-side gate-drive circuit is located right on top of the ground plane for the DC bus. Even though the capacitance may be relatively small, the ΔV is large – the full bus voltage. Moreover, the dv/dt across this capacitance is the fast dv/dt of the switch-node. The simplified parasitic PCB capacitance is drawn in red on the schematic in **Figure 15**. This type of layout will most likely experience problems due to the charge injected into the gate-drive circuit on every switching edge.

This is an example of one of the compromises or tradeoffs that often must be made during a PCB layout: optimizing the gate-drive layout results in a less-optimal power loop. Conversely, optimizing the power loop produces gate-drive layout problems. There is another approach to the gate-drive layout to consider: keep the double-sided return path for the power loop and move the gate-drive circuit to the back-side of the PCB. The gate-drive loop would then need to include vias to connect to the transistor. Side-by-side vias provide a reasonably low impedance, and the back-side of the board is farther away from the ground plane. This seems like a promising approach, but there is still another layer to work with that can provide a unique solution.

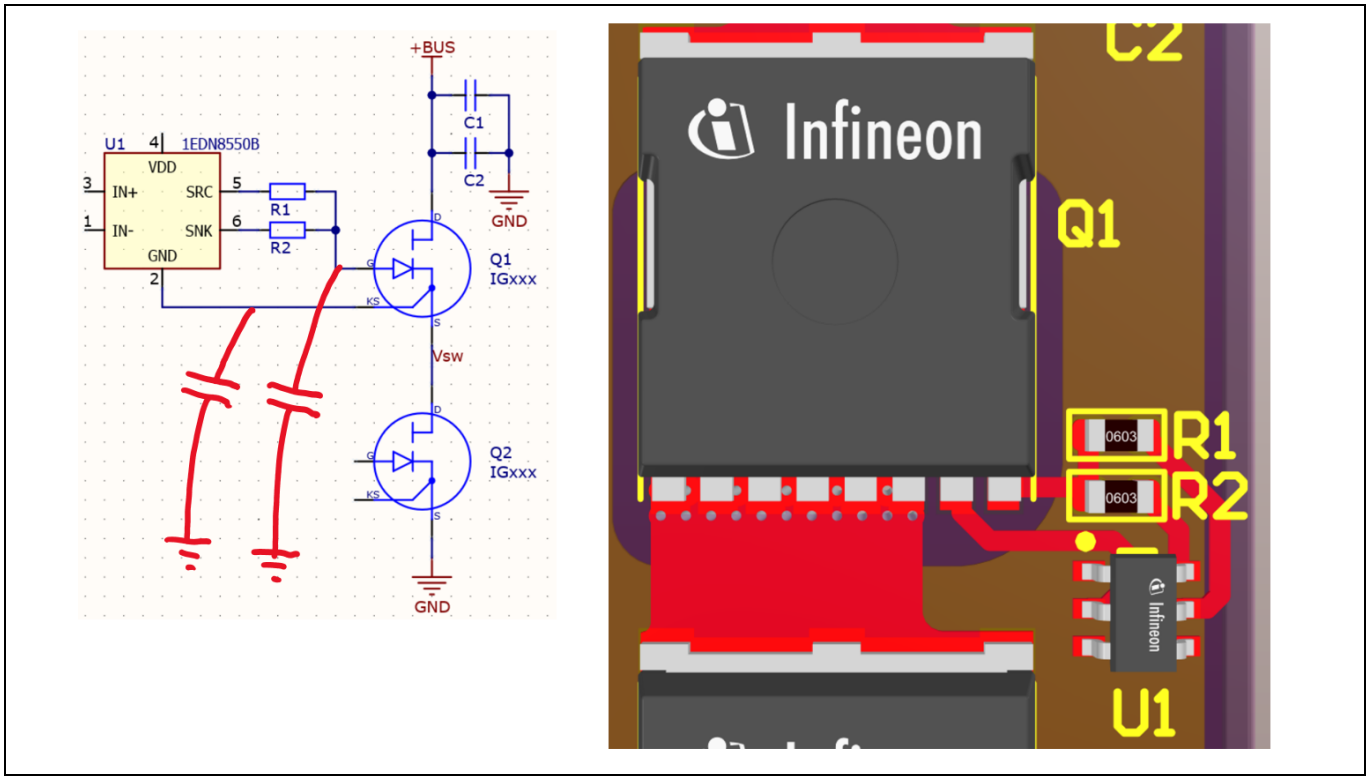

Figure 15 **Double-sided return path interferes with gate-drive routing and adds common-mode capacitance – not recommended**

9 Using a driven Faraday shield

With the gate-drive circuit on the back-side of the PCB, layer 3 is available as a return plane. The goal here is to minimize the effect of the common-mode capacitance between the gate-drive circuit and the bus ground plane. If a local plane is added on layer 3 (cyan color), and connected to the switch-node, the value of the capacitance changes, and also the current path to charge and discharge that capacitance.

The switch-node is the “local ground” for the high-side. Adding the plane on layer 3 does not really differ from the return plane being added to the gate-drive circuit in **Figure 13**; the capacitance between the circuit and the plane will be the same. However, the layer 3 plane shields the gate-drive circuit from the problematic capacitance to bus ground. The capacitance to bus ground doesn’t go away, but it is now between two low-impedance planes rather than coupling to the gate-drive circuit and its components. Thus, the gate-drive circuit only “sees” capacitance to its own local common, not to the bus ground plane. The current necessary to charge and discharge the inter-plane capacitance therefore comes directly from the low-impedance switch-node, completely bypassing the gate-drive circuit **[2]**.

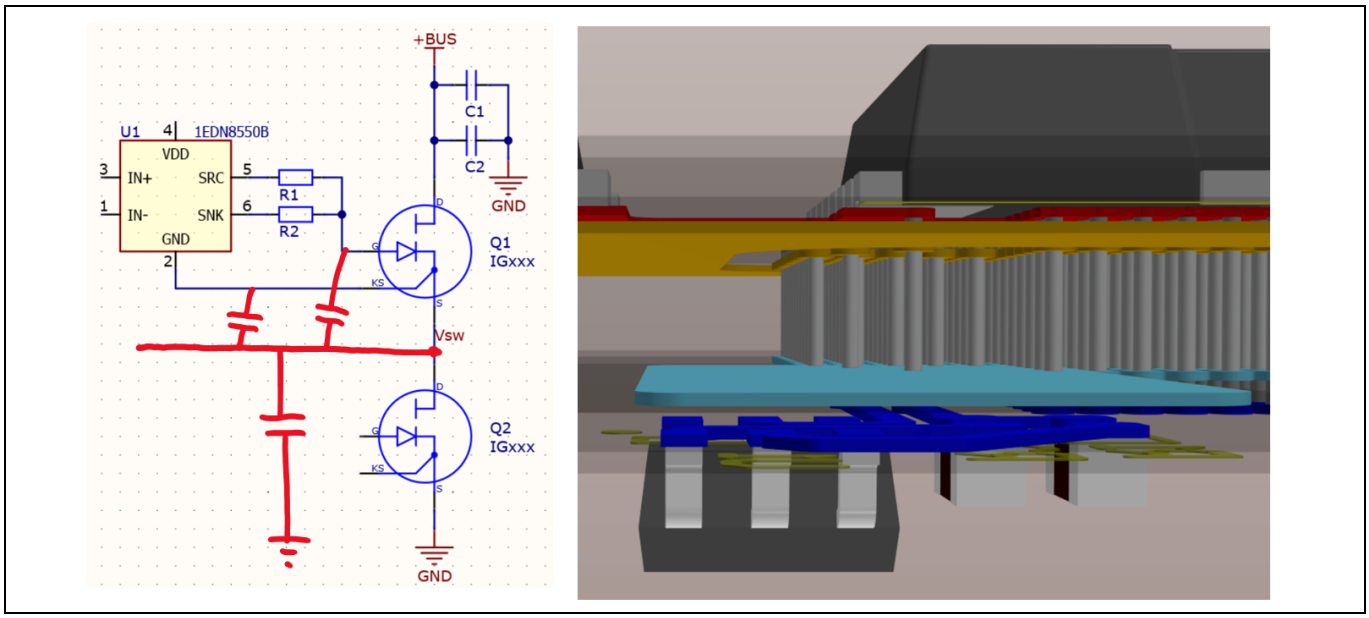

Figure 16 Driven Faraday shield to mitigate the common-mode capacitance problem

Avoid driving the Faraday shield from the actual gate-driver ground – the Kelvin source (KS), rather than the switch-node. While the small voltage bounce between S and KS is insignificant compared to the other problem this connection would cause, by forcing all of the transient charge/discharge current to flow through the KS bondwire(s). The induced reaction voltage $v = L\,di/dt$ appears as a differential gate-drive signal and can be a problem not only for shield capacitances, but even for “unintentional” capacitances, like the capacitance across any component connecting the high-side circuit to bus ground.

Figure 17 depicts a typical example. The high-side gate-drive circuit is capacitively coupled to the low-side ground through a signal isolator, and a DC-DC transformer (or alternatively the junction capacitance of a bootstrap diode). While these are typically only a few pF, there are examples of DC-DC converters for gate drive that have more than 50 pF of capacitance across the isolation barrier. As shown on the right of **Figure 17**, when the half-bridge is capable of switching 400 V in a few ns, every pF is important. At 100 V/ns switching slew rate, each pF results in 100 mA peak common-mode current. The injected current spike then has to return to bus ground through the low-side circuits, and along that path, it could cause glitches in the logic or other circuits.

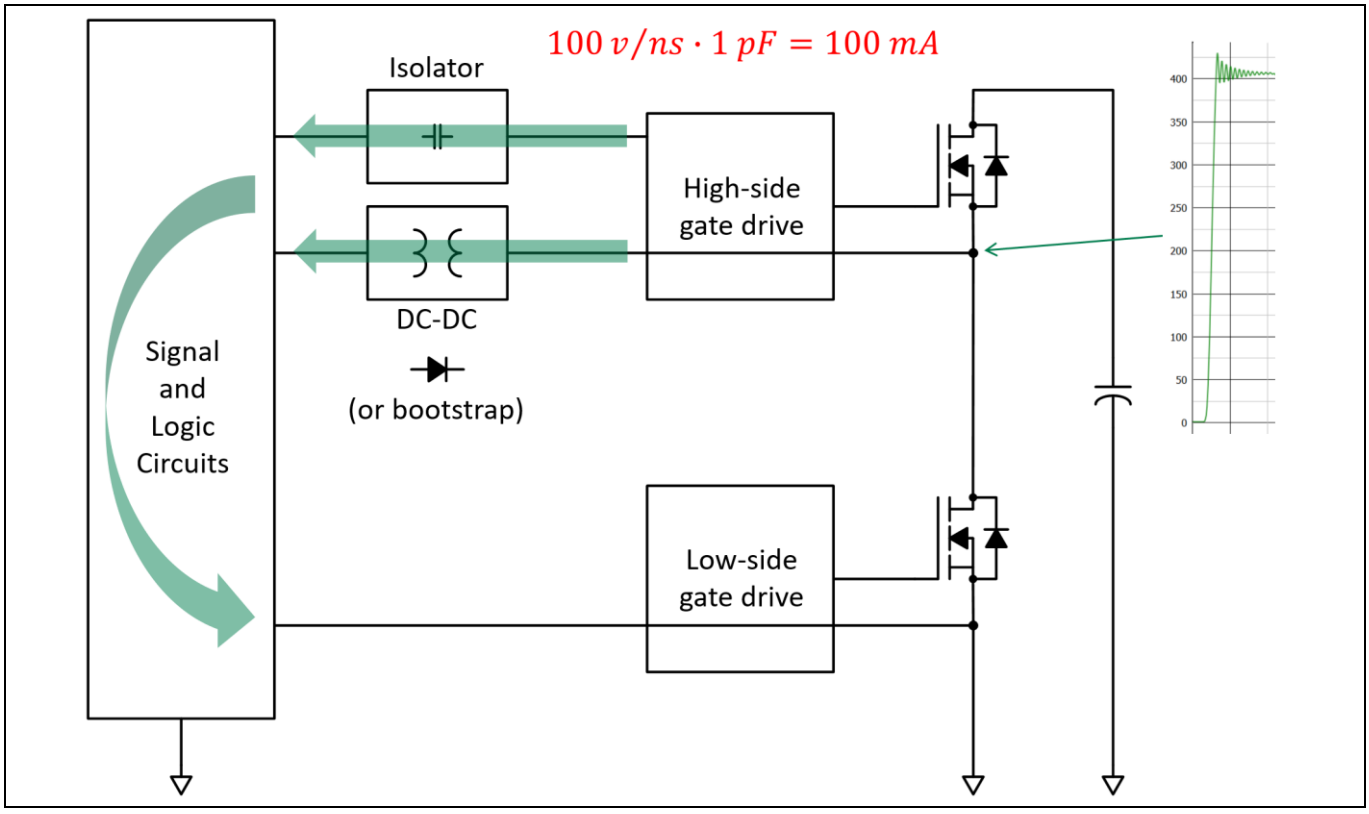

Figure 17 "Unintentional" pathways for common-mode current

There is really no way to completely eliminate this capacitance, so the best approach is to think about where the return current will flow, and to minimize any effect – make sure the return current has a low-impedance path back to bus ground. Note that these capacitive effects can be just as much a problem for integrated driver + transistor as a discrete design – it all depends on the interconnection inside the integrated driver + transistor.

To better understand the problem caused by the parasitic capacitance charge/discharge current, consider the system shown in **Figure 18**. The switch-node is the low-impedance driving point (driving the fast-switching transition), shown in red. In addition, everything in orange is also connected to this node. The difference is that everything in orange is driven *indirectly* – through the KS pin of the transistor. The KS pin is only intended to be a part of the differential gate-drive loop, and it is commonly just a single bondwire like the gate. Due to the proximity inside the transistor package, G and KS bondwires have low loop area and mutual inductance helping to minimize overall gate-loop inductance.

However, when the current through the KS bondwire is NOT common to the gate bondwire, there is no mutual inductance cancelation and therefore a differential-mode voltage will be generated within the gate-loop. As the switch-node voltage rises, the voltage drop across the KS bondwire subtracts from the applied gate-drive voltage, and thus slows down the switching transition. This effect can essentially cancel the benefit of the Kelvin source package by slowing switching speed much in the same way that common-source inductance slows switching speed in a conventional three-pin transistor package.

An additional concern illustrated in **Figure 18** is spreading out the switch-node. It is best to keep the switch-node compact, to 1) keep the capacitance minimized, and 2) prevent it from radiating and capacitively coupling to other parts of the circuit. Rather than placing the gate-drive DC-DC converter on a separate board, it is best to put that function right next to the power stage – on the same PCB as the half-bridge.

The situation in **Figure 18** should be avoided. Routing the orange traces all over the system causes distributed capacitive coupling all along the path, and it radiates onto several PCB paths. On the aux supply card, there is

the additional capacitance across the transformer isolation – so that injected current will then travel back to ground through the main PCB – potentially causing noise and interference along the way.

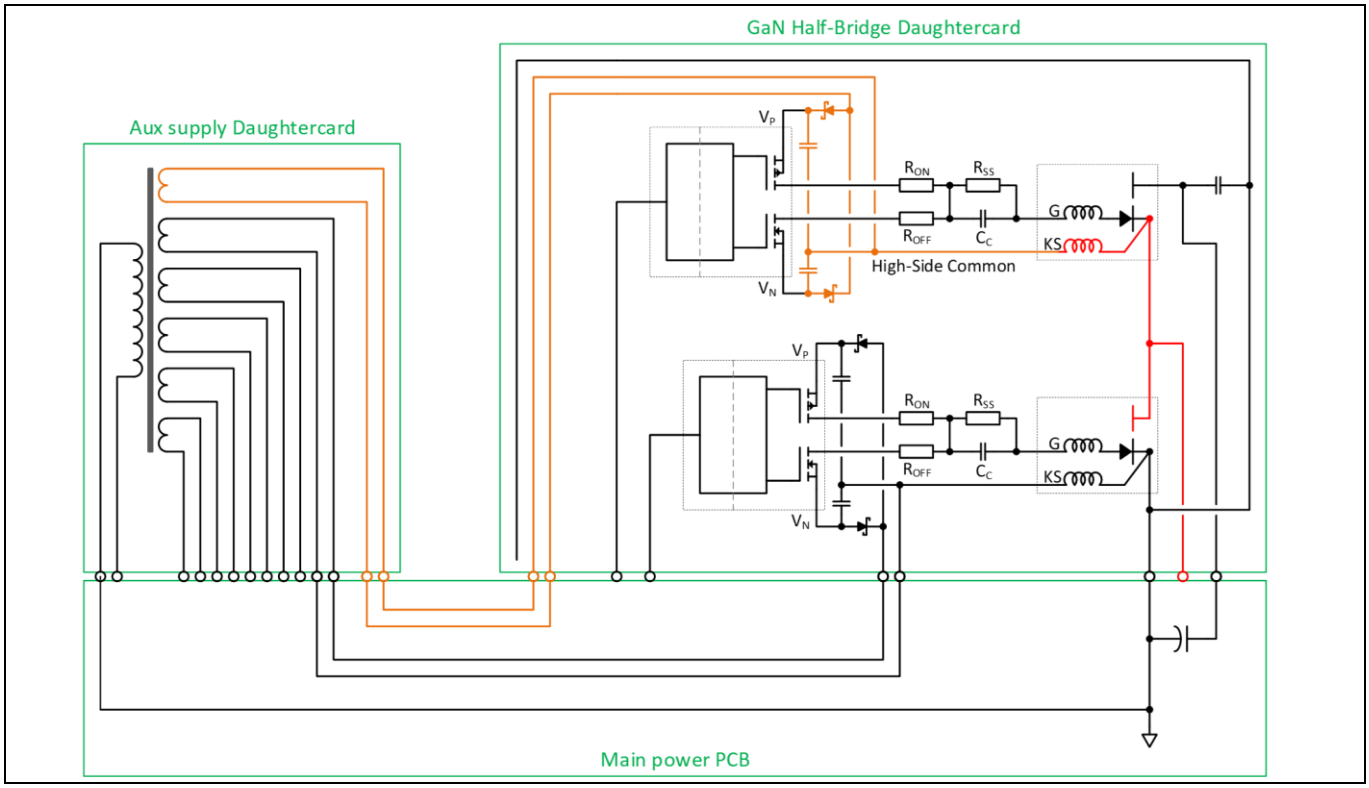

Figure 18 Common-mode currents can cause differential-mode gate-loop voltage

An example of a good GaN half-bridge layout is shown in **Figure 19**. This is a half-bridge "daughtercard" that includes the two GaN transistors, high-frequency bus capacitors, isolated high- and low-side gate drivers, and the isolated DC-DC supply for both high- and low-side. A single heatsink is attached through a thermal interface material by a spring-clip (not shown). This half-bridge circuit is the high-frequency leg of an active front-end rectifier (bi-directional PFC), which can deliver 3.4 kW continuously to/from 240 V AC-line.

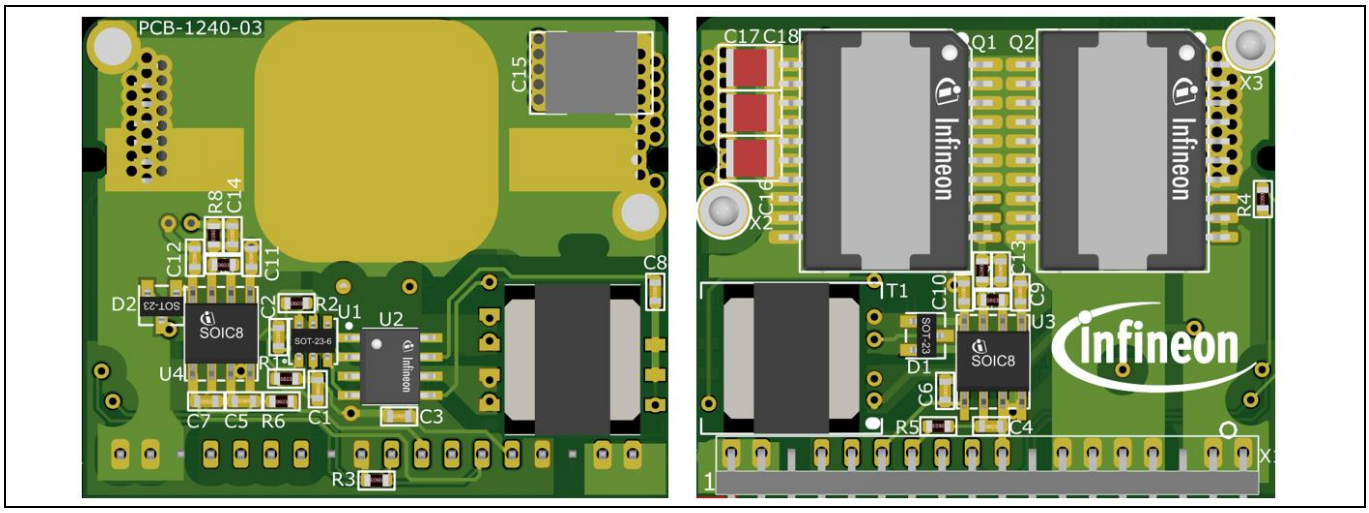

Figure 19 Example of a 3.4 kW complete GaN half-bridge measuring 32 x 39 mm

10 Summary of the key recommendations for optimizing performance of high-speed HV GaN transistors

The table below provides a summary of the key points discussed in this article.

Table 1 Recommendations and further considerations for optimizing performace of high-speed HV GaN transistors

Recommendation	Further considerations
Consider where current will flow during switching transitions	Be sure to include parasitic elements in that assessment. Remember to include the complete return-path of the curremt in the analysis.
Layout inductance may be critical in some parts of the circuit, but unimportant in others	Recall the “half-bridge as a two-port network” section
Minimize layout inductance by taking advantage of PCB layer pairs with thin dielectric	Route outbound and return currents along the same path, but on adjacent parallel layers in opposite directions
Avoid deviations from the “over/under same path” that will result in lateral loops	As well as lower mutual inductance and therefore higher loop inductance
Remember: package inductance is not necessarily a fixed value for any SMT package	The value should be dependent on the assumed return path
Use top-side cooled SMT packages to optimize both the electrical and thermal paths, independently	Without having to make compromises or increase cost from adding more than 100 thermal vias
Use a plate for the return-path of gate-drive circuits	Directly below the circuit, connected to the KS pin
Prevent capacitive currents	From switch-node to ground from flowing through KS pin
Keep ground reference circuits away from high-side gate-drive circuit	Or use a driven Faraday shield if that is not possible
Keep the switch-node compact	To minimize capacitance and radiation into other parts of the circuit

The fast-switching capability of GaN transistors can make PCB layout more challenging. This article discussed several key concepts to help the user understand the layout challenges, and strategies to help solve these challenges and optimize the layout for best overall electrical and thermal performance.

Following the recommendations explained in this document and summarized in the table above will help designers obtain optimal performance from high-performance GaN technology.

References

[1] Clayton, Paul: *Inductance: Loop and Partial*; Wiley; 2010
[2] Persson, Eric: *PCB Layout Techniques for Optimizing Performance of Surface-Mounted Wide-Bandgap Power Electronic Circuits*; IEEE APEC 2022; Seminar S02

Revision history

Document revision	Date	Description of changes
V 1.0	2023-07-28	Initial release

Trademarks
All referenced product or service names and trademarks are the property of their respective owners.

Edition 2023-07-28

Published by

Infineon Technologies AG

81726 Munich, Germany

Do you have a question about this document?

Email: erratum@infineon.com

Document reference

AN_2306_PL52_2308_151809

Important notice

The information contained in this application note is given as a hint for the implementation of the product only and shall in no event be regarded as a description or warranty of a certain functionality, condition or quality of the product. Before implementation of the product, the recipient of this application note must verify any function and other technical information given herein in the real application. Infineon Technologies hereby disclaims any and all warranties and liabilities of any kind (including without limitation warranties of non-infringement of intellectual property rights of any third party) with respect to any and all information given in this application note.

The data contained in this document is exclusively intended for technically trained staff. It is the responsibility of customer's technical departments to evaluate the suitability of the product for the intended application and the completeness of the product information given in this document with respect to such application.

Warnings

Due to technical requirements products may contain dangerous substances. For information on the types in question please contact your nearest Infineon Technologies office.

Except as otherwise explicitly approved by Infineon Technologies in a written document signed by authorized representatives of Infineon Technologies, Infineon Technologies' products may not be used in any applications where a failure of the product or any consequences of the use thereof can reasonably be expected to result in personal injury.