Addressing EMC Preemptively in Ethernet Connected Devices

written by Fred Greenfeld

Introduction

Designing an EMC compliant Ethernet-connected device can be a frustrating experience. EMC compliance testing typically occurs late in the schedule when the product is in verification testing. A failure at this point can dramatically impact the development schedule.

Of course, the engineer will have taken steps to address EMC based on previous design experience and simulation exercises, but such effort does not guarantee an EMI compliant design. A new product design is by definition “new”, and therefore different from previous designs. There may be changes that affect EMC performance in unexpected ways. What worked before may not work well enough this time.

Simulation results are only as good as the simulation model used. Unless parasitic elements are considered, including PCB parasitics, the results may be misleading. However, accurately determining the parasitic elements of the PCB requires the PCB layout be completed. Such a strategy is not consistent with pre-emptive EMC mitigation since results are delayed until after the majority of the design work has already been completed. If a problem is discovered, the solution takes longer to execute and has a greater impact on the schedule.

So, how does the design engineer design an EMC compliant device without complete information on the EMI sources and the parasitic coupling paths?

The discussion which follows is not comprehensive. It cannot address solutions for every situation, for such is beyond the scope of the document. The intent is to encourage the designer to consider the “big picture” before delving into the details.

Preemptive Design

First and foremost, every design methodology strives to achieve 100% first pass design success. Despite the development team’s tools, resources, and design experience, success in achieving this goal is typically less than ideal. There are myriad reasons as to why this happens including compressed schedules, tools availability, and stretched resources. However, even in the absence of these reasons, 100% first pass design success rarely occurs. Why is this? The simple answer is that it just takes too long to model a system in sufficient detail to predict every significant aspect of its operation. The philosophy will always be do the best job possible with the resources available in the allotted time, build evaluation units, discover where the problems are and take corrective action for the next iteration. As tools and resources improve the situation will improve, but how long will this take and what can be done in the shorter term? What other strategies may be employed?

The key is to create a flexible design so that there are options available that do not require much, if any, redesign effort. In other words, do not paint yourself into a corner. Have a back-up plan. Of course, it is not possible to anticipate where every problem will occur, and it would be easy to take this strategy too far and over-design the product, but a considered effort can be made based on history and experience.

The IC design community has followed this strategy for a very long time. Every aspect of the die layout is considered to maximize silicon utilization. Routing and device placement are deliberate, considered actions. Spare devices are packed into the blank areas of the die just in case they are needed. During wafer fabrication wafers are held at intermediate processing steps so that interconnect revisions can be made without the delay of a full fabrication cycle. The re-design cycle time is greatly reduced.

Although not completely applicable to PCB design, a strategy like the one employed by the IC design community could be used. How many times has a product design been compromised by PCB component placement and trace routing? How many times has the PCB layout caused signal coupling or noise injection that resulted in adding filtering components and/or a PCB design revision? The PCB designer, whether an engineer or not, may not have the insight or experience to make the optimum layout choice. Sometimes schedule pressure results in rushed layout decisions. Either way, if component placement and routing order priority is not deliberately set and followed, the PCB layout will have an element of arbitrariness which can compromise product performance. Furthermore, the placement of component spares, as is done by the IC design community, requires careful consideration and a different strategy. Component size restricts where spares may be placed and routing traces to them may compromise their utility. More thought is required to place them near where they may be required, and if they are not needed they may be omitted or replaced with 0-ohm jumpers as appropriate. The placement of spares should occur after initial component placement and critical trace routing. Once initial placement and routing is done, areas of concern may be identified. Obviously, there is a practical limit to what can be done with spares and pre-emptive design without adding to the problem, so let's look at some practical examples.

PWB Layout

The first step in the preemptive design philosophy is to get the PCB design done well. The following discussion does not present rigorous background theory. The intent is to suggest areas for the designer to consider during PCB design. If rigorous theory is desired the references listed at the end may be consulted.

The impact of the PWB layout to the success of the product design is as critical as any other component, and it often does not receive the design scrutiny required. This is particularly true for EMI mitigation in high speed digital signals and analog signals with high di/dt or dv/dt such as occur in power circuits. So, before we examine design flexibility options, it is important to review fundamental layout guidelines for PCB design as it pertains to EMI and signal integrity.

At the most basic level the source and return path of any signal forms a circuit which creates a loop antenna. The strength of the radiated signal is proportional to the loop area, the current flowing through it, the length of the signal path, the frequency of the signal, and the impedances of the source and loop. Furthermore, the direction of the radiated signal depends on the loop path length compared to the wavelength of the signal. For wavelengths on the order of the length of the loop the radiated signal is perpendicular to the plane of the loop. For wavelengths much longer than the loop length the radiated pattern shifts to the plane of the

loop. The wide frequency spectrum of fast digital signals means the radiation pattern will vary widely. All this aside, the point is to minimize both the loop length and area to minimize radiated emissions. If it is not possible to create an acceptable loop length and area then the alternative is to shield it. It should also be noted incomplete shielding can result in the creation of a ground plane reflector which may enhance radiated emissions in the unshielded direction. While this is typically not an issue for the layer-to-layer separation in a typical PCB since the distances of concern between the loop and the reflector are on the order of 0.05 to 0.25λ, the distance to the side of a grounded enclosure should not be overlooked as it may amplify emissions directionally. For example, consider a 1000Base-T (1gigabit) Ethernet application. It requires a minimum 100Mhz BW for each of the four twisted pairs of the Ethernet cable. The wavelength of 100MHz is 3 meters, and 0.05*3 m = 150 mm. This is a large distance compared to the layer spacing of a PCB, but perhaps not so much for the distance to the side of an enclosure. Now consider higher data rate Ethernet application which can require up to 2000MHz BW (25Base-T and 40Base-T). The distances of concern are reduced to 7.5mm.

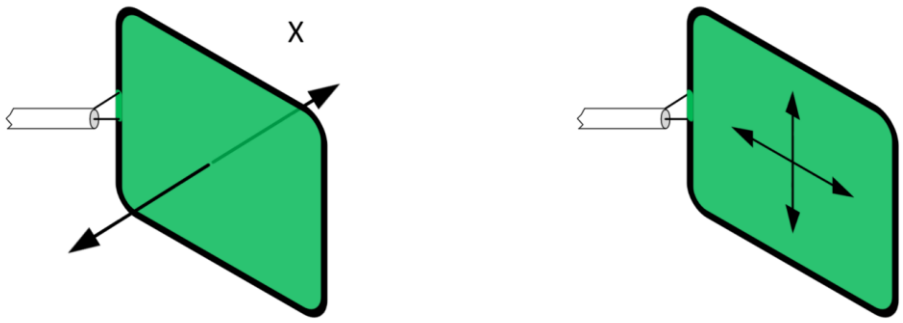

Figure 1. Loop Antenna Radiation Direction: (a) Loop Length ≈ λ, (b) Loop Length < λ

Other loop antenna paths are formed in the common mode (CM) signal path with Earth ground. All AC signals capacitively couple to their surroundings. Even though the capacitance is small and the currents tiny, the loops can be very large. Current flows through the parasitic capacitance and finds a return path which forms the loop antenna. Since it is impossible to eliminate the parasitic capacitance, the strategy is to make the path length short. The interwinding capacitance of an isolation transformer, for example, allows current flow between the primary and secondary circuits. The return path may be difficult to identify, but there will be one. The path may be through Earth ground, a conductive enclosure, or other parasitic connections. The only certainty is the loop will be large. To create the smallest loop possible, appropriately rated capacitors can be placed between primary and secondary circuit grounds as close as possible to the transformer, and with clever PCB design, capacitance can be created using overlapping primary and secondary ground planes on separate layers of the PCB.

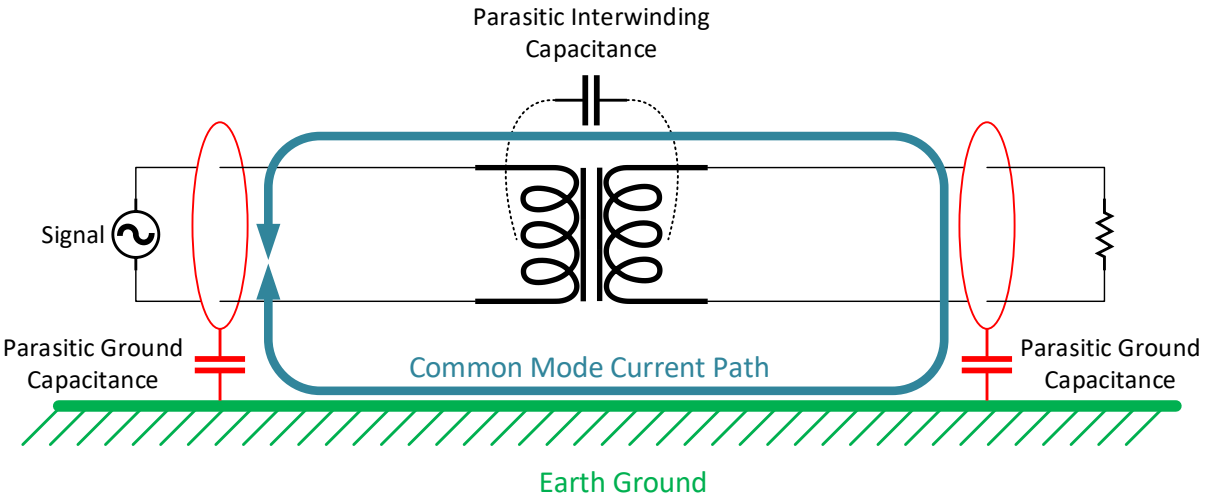

Figure 2. Common Mode Current Path. Secondary side signal ground may be directly connected to Earth Ground.

Reduction of loop size is the reason systems connected to Earth ground typically have signal ground connected directly to Earth ground. For systems requiring source isolation (AC mains powered applications, for example) safety standards do not allow primary ground to be directly connected to Earth ground. While the isolated (secondary side) signal ground may be directly connected to Earth ground to minimize the common mode circuit path, the primary side signal ground must remain isolated. Part of the solution is to couple the grounds capacitively using safety agency approved capacitors to provide a shorter path, but this only provides a shunt path and does not eliminate parallel paths. The common mode currents will divide among the alternate paths in proportion to the path impedance like a parallel network of unequal resistor values. Most of the current will flow through the intended path (assuming its impedance is significantly less than the alternatives), but not all. To block the undesired current paths a series impedance must be added, usually in the form of a common mode inductor. An obvious example is the conducted CM EMI filter found in power supplies. The CM inductor impedes CM current flow toward the power source and the capacitors to ground provide a shunt return path.

Another common mode current source to contend with is differential to common mode conversion in unbalanced differential pairs. Ethernet signals are transmitted over twisted pairs in the cable and converted to a microstrip or stripline transmission line configuration on the PCB. If the series and shunt impedances are not identical in both signal paths, a common mode current will flow. The impedance mismatch may be caused by unequal parasitic capacitance to ground or even minor differences in the signal path lengths.

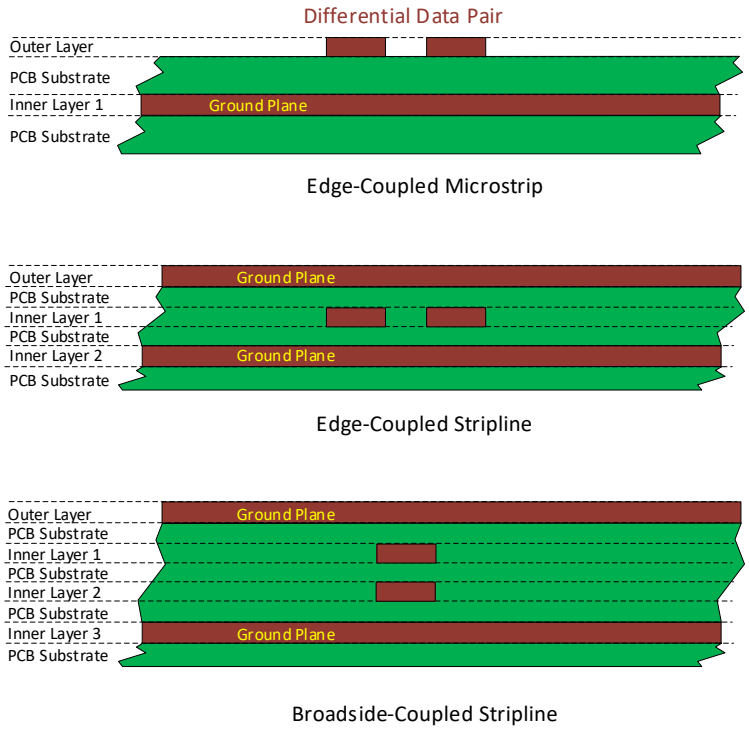

Figure 3. Common PCB Transmission Line Configurations. Each configuration has different impedance characteristics.

The principal signal integrity parameters affected by the PCB design are insertion loss, return loss, and crosstalk. In simple terms, the insertion loss measures the signal power attenuation between source and load. Return loss measures the portion of the transmitted signal that is reflected back to the source. Crosstalk measures the coupling of adjacent signals, differential pairs in this case. The key take-away is to maintain the characteristic impedance on the PCB using a low-loss transmission line configuration and maintain adequate separation between adjacent differential pairs. Ethernet cables (CAT5, CAT5e, CAT6) have a characteristic differential impedance of 100 ohms on each twisted pair with a real series resistance of < 0.188 ohms/m. When the signal transitions to the PCB, the characteristic impedance should be maintained to minimize the impact on emissions and to prevent degradation of signal integrity.

Designing the signal traces to a characteristic impedance is well documented so the theory and methodology shall not be repeated here. PCB impedance calculators are readily available online or accessory to layout tools. However, achieving the desired performance criteria may be a different matter entirely. An ideal layout would route the differential signal from the source to the load using the shortest equidistant straight-line path on one PCB layer. While single discontinuities (sharp corners, vias, changes in ground plane coupling) in the signal path may be ignored, their cumulative impact may be significant especially with broad spectrum applications.

A sharp trace corner creates shunt capacitance to the ground plane and increases the local E-field strength. The former degrades insertion loss while the latter increases radiated

emissions. If trace directional changes are required, they should be smooth and continuous with a radius no tighter than the differential pair separation. If the layout tool does not support smooth curves, then smooth curves can be approximated with cumulative turns of no more than 45°. If a sharper angle is required, such as a right-angle, the outside corner should be rounded or chamfered to minimize the capacitance and reduce E-field strength.

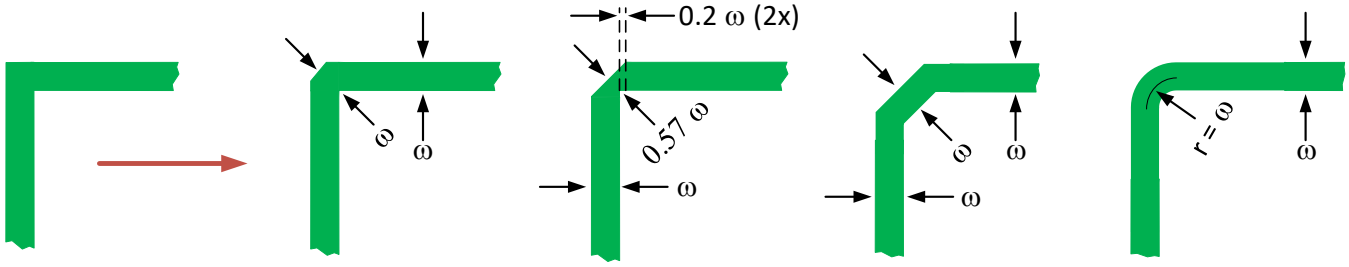

Figure 4. Chamfered corner options

Vias introduce discontinuities which impact the characteristic impedance. Any discontinuity will degrade insertion and return loss. While vias create both inductance and capacitance, the inductance has the greater impact on signal integrity. The inductance is proportional to the length and diameter of the via with length being the dominant contributor. As a guideline, each via in a typical 0.063-inch-thick PCB adds ~ 1nH of inductance.

Capacitance is created between the annular rings surrounding the via and internal ground plane(s). This includes the annular ring on the same layer as the ground plane as well as outer layer annular rings, particularly if the outer rings are larger than the inner ones. Like any plate capacitor, the capacitance may be reduced by increasing the separation between plates and/or reducing the area of the plates. Therefore, reduce the diameter of the annular rings as much as possible while increasing the diameter of the plane void (anti-pad). While the capacitance of a typical via is very low, on the order of 0.5pF (30 mil annular ring with 10 mil spacing to the plane), it is cumulative with each additional via.

If vias cannot be avoided care must be taken to maintain the characteristic impedance when the signal path transitions to another PCB layer. For differential signals routed on an outer layer the signal traces are routed in a micro-strip configuration having a single ground plane adjacent to the differential pair. Once the signals transition to an inner layer the ground plane geometry will be significantly altered. There may be a plane above and/or below, broken planes interspersed with signal traces, or just signal routing on adjacent layers. Transitioning from one transmission line configuration to another requires dimensional adjustments to maintain the characteristic impedance. The characteristic impedance changes depending on what is routed on adjacent layers. Characteristic impedance depends on the geometry of the traces in relation to each other, to the plane(s), and to nearby signal traces. The separation between the differential pairs, the distance to the plane, the trace width, and trace thickness all impact characteristic impedance. Since trace thickness depends on the PCB layer copper weight, characteristic impedance is impacted if the copper weight changes. Many PCBs, particularly applications which embed DC-DC converters, specify heavier copper weights on exterior layers. While it is possible to compensate for the changes in characteristic

impedances, it is much simpler to avoid the problem altogether if the routing of the signals is prioritized so that layer changes are unnecessary.

For single-ended signals a PCB design issue often overlooked is signal return paths. Most designers create interconnected ground planes and assume they have created the lowest impedance return path, and this would be the case if there were no vias or signal routing on ground plane layers to interrupt the current flow. However, this is seldom the case and is often not revisited once routing has been completed. The return current follows the lowest impedance path back to the source. For DC and low frequency signals the lowest impedance path is dominated by resistance and current follows the shortest distance. For higher frequencies the impedance is dominated by inductance. As frequency increases the lowest impedance path moves underneath the source path where inductance is at a minimum due to counterflowing currents. It is important the return path is not interrupted by via clusters or signal traces in the ground plane. To the extent the return path deviates from the ideal there will be degradation to signal integrity due to increased impedance. Since the impedance increase is predominantly inductive, high frequency noise will increase because of signal di/dt. The signal edges will become noisy and degrade circuit performance.

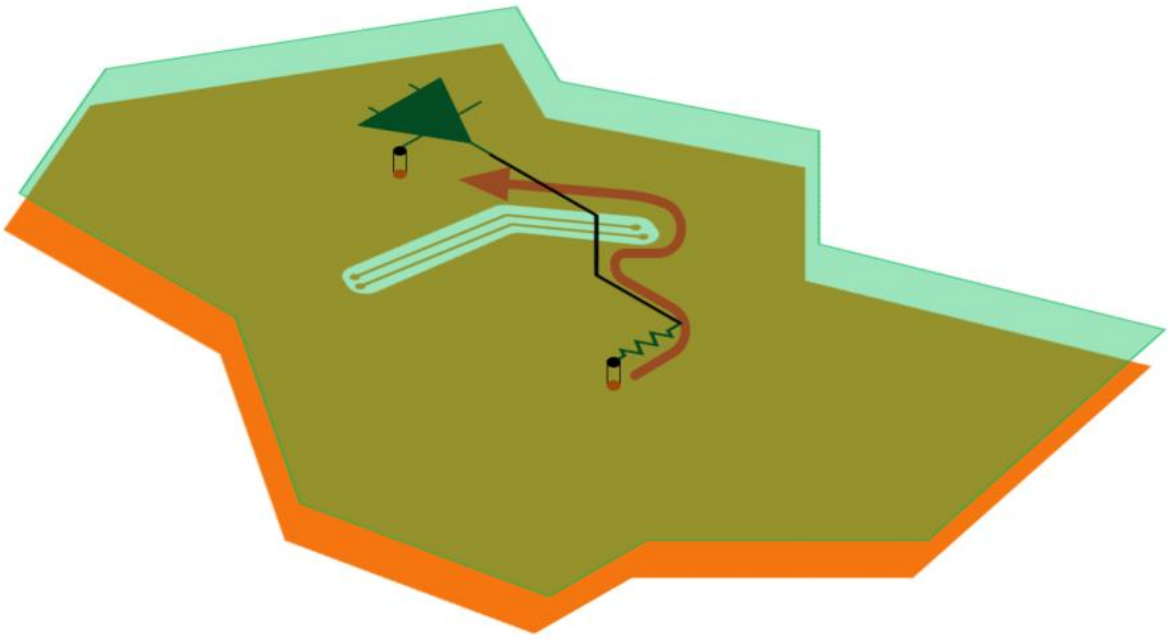

Figure 5. Discontinuous ground plane forces deviation of return current path

However, there are situations where it is advantageous to split planes into sections. The number of PCB layers available for trace routing planes is limited by cost considerations. The result is that mixed signal types will share a common plane. This is the case when there are mixed signal types such as high di/dt or dv/dt signals utilizing the same ground plane as high impedance signals. It is already common practice to separate grounds into digital, analog, power, and others, but even within the same category it is often necessary to channel currents to minimize the impact of parasitic elements on adjacent circuitry. This is the strategy behind

star ground connections but can be taken a step further when a ground plane layer is available. A star ground consists of a single ground reference with separate traces radiating outward to the respective circuitry. While currents are kept separate, the ground path itself has impedance dependent on the length and width of the trace. Long or narrow ground traces may not provide a sufficiently low impedance. In contrast, the partitioned ground plane, having a centralized ground reference (using several closely spaced vias) where the sections are minimally separated by narrow voids (anti-traces), provides both low impedance and segregated currents.

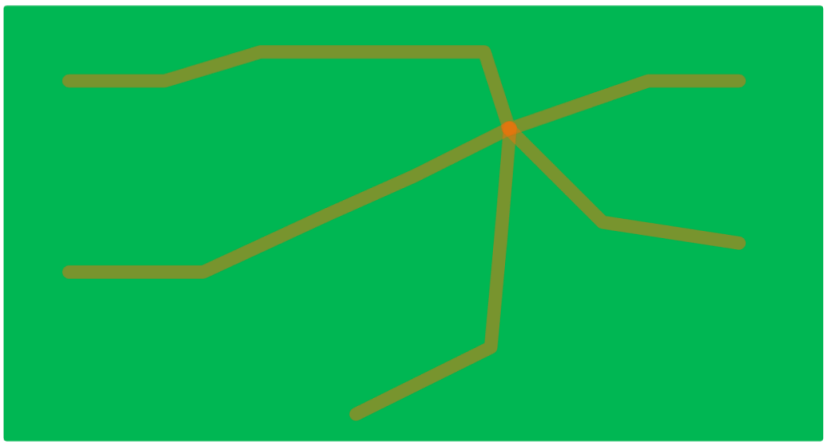

Figure 6. Star Ground

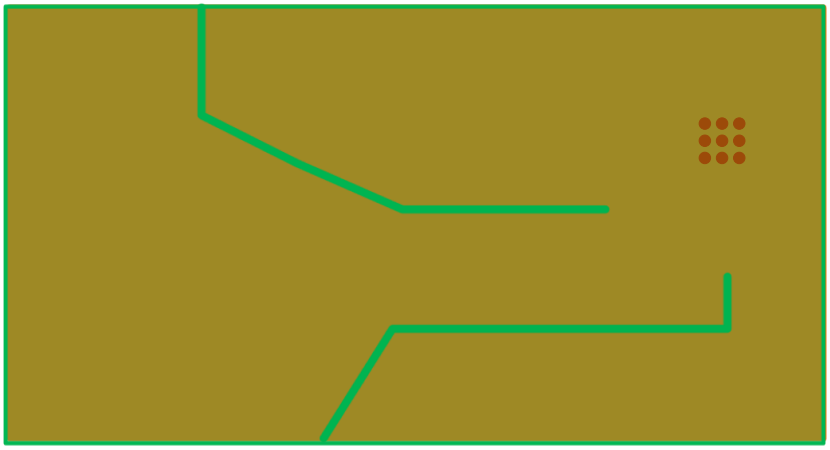

Figure 7. Partitioned Ground Plane

Confining the current path is also required when using capacitive filters. Capacitive filters are most effective when the signal path connects directly to the capacitor lead rather than through an indirect “T” connection. To the extent an indirect connection is made, parasitic resistance and inductance are introduced which add to the capacitor’s intrinsic equivalent series resistance (ESR) and equivalent series inductance (ESL) resulting in less effective filtering. The basic rule is “route the signal to the capacitor rather than route the capacitor to the signal.”

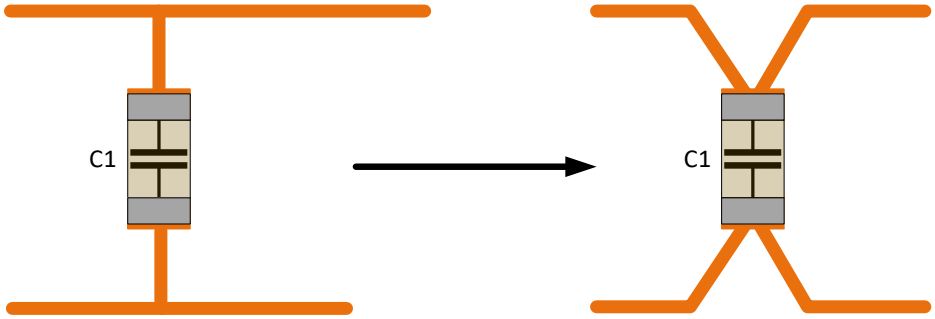

Figure 8. Improved capacitor routing

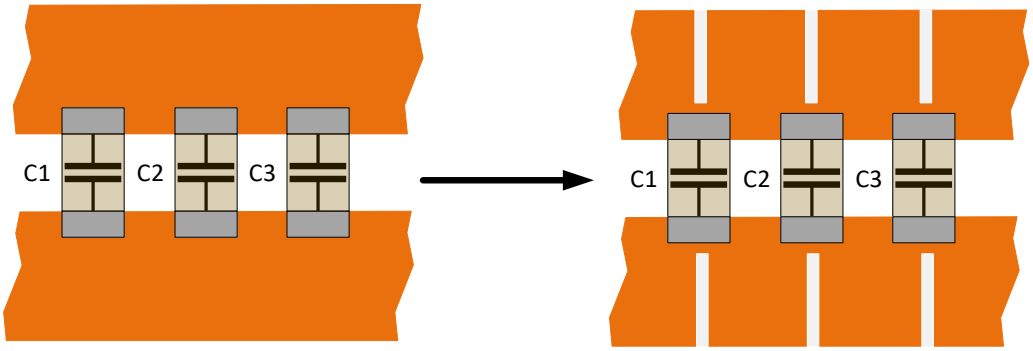

Figure 9. Improved capacitor routing for high currents

Signal integrity is also compromised by crosstalk, the undesired coupling of one signal to other signals through inductive and/or capacitive coupling, antenna coupling from a radiated source, and ground shift due to impedance in a shared return path. Understanding the coupling mechanism is key to minimizing crosstalk.

Antenna coupling was previously addressed, although from the perspective of a source. The same methods may be used to reduce the effects on a target signal from a radiating source, but it is usually a better strategy to go after the source since there are likely more targets than sources. If the effective loop size and area cannot be adequately minimized, then shielding may be used. Some applications, such as those that deliberately generate RF signals, partitioning and shielding of sensitive circuit blocks is a necessity.

Capacitive coupling is a function of the signal trace separation and the capacitor plate area. For signals that cross on separate PCB layers, the plate area is dominated by the overlap region. However, capacitive coupling also occurs with signals sharing a parallel path where the paralleled traces form the capacitor plates. Furthermore, parallel signals on the same or adjacent layers form a transformer structure which will couple inductively. Minimize the overlap area and the length of parallel traces by crossing at a 90-degree angle, use the minimum acceptable trace widths as constrained by characteristic impedance and current where they cross, and separate the signals to the extent practical.

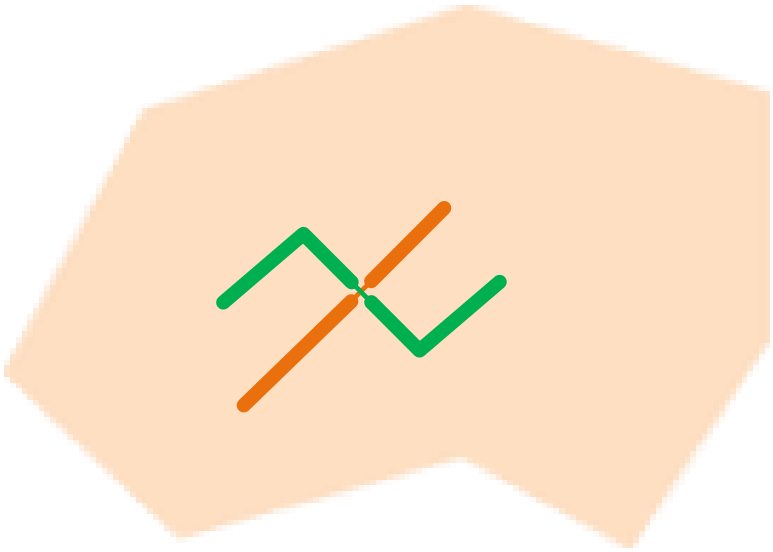

Figure 10. Minimize trace crossing area and separate parallel traces

Crosstalk may be further minimized by using guard rings. Guard rings are low impedance ground traces placed around or between the signals of interest. The concept of placing grounds above and below signal layers (i.e., ground planes) has already been discussed. Expanding on that concept, we can place grounds adjacent to signals on the same layer. By inserting a ground trace between a source and target, not only is the signal separation increased which reduces coupling, but it creates a three-terminal capacitor with the middle plate at a fixed ground potential. Instead coupling to an adjacent signal, the coupling is to ground. The capacitive loading of the signal is not improved, but it is no longer coupled to nearby signals. If ground is not available, any DC potential may be used so long as it is low impedance.

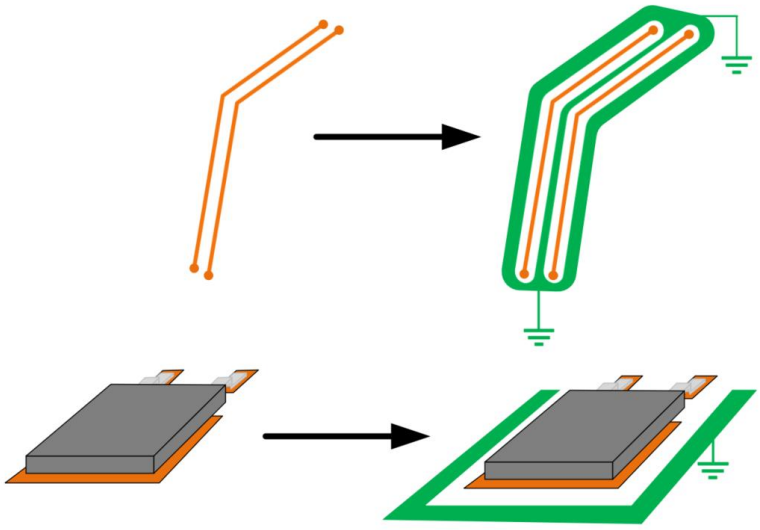

Figure 11. Introducing guard rings to shield adjacent circuitry

The same technique may be used to reduce inductive coupling. As stated earlier, long parallel traces will inductively couple. A guard ring separating the signals works in a similar way as a twisted pair by providing a counter flowing current path. In this case it is for the induced currents rather than the return currents. Like the capacitive coupling situation, signal separation is increased when a guard ring is introduced, but the majority improvement comes about due to the magnetic field cancellation caused by the counter flowing currents in the guard ring. The only caveat is the current path through the guard ring must be a closed loop. If the loop is not closed there is no induced current, although there will be induced voltage just as there would be to an adjacent signal. The easiest way to get a closed loop is to connect both ends to a ground plane. Additional intermediate connections along the trace length reduces impedance.

Switching nodes, particularly those found in the power supply, should be shielded by a ground plane to minimize radiated emissions. The capacitive coupling to the plane, or adjacent circuitry if not shielded, results in current spike injection from the dv/dt of the switching node. If the return path is long or of high impedance ground shift will result. These currents should be contained locally to avoid coupling to other signals. The easiest way to deal with these currents is to shield the switching node with a separate local ground plane in the outline of the switching node, like a shadow. The injected currents are now confined to the local plane. To reduce the di/dt of the injected current the plane should complete the return current path through a resistor. In effect an RC snubber is formed which limits both di/dt and dv/dt. As an example, consider an N-channel switching FET. The drain and connected components form the switching node which capacitively couples to the adjacent layer(s). The copper trace pattern of the switching node is copied into the adjacent layer(s) and connected to the FET source terminal through a resistor.

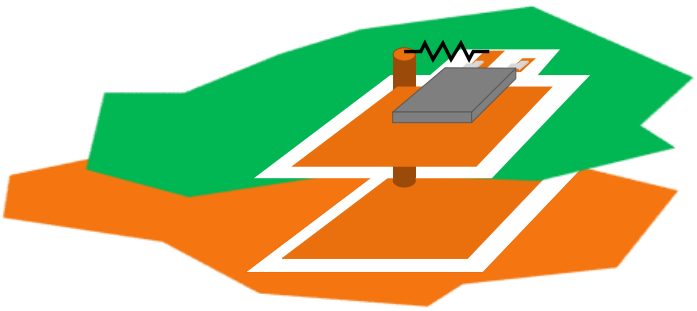

Figure 12. Resistor Connected Shield

Perhaps the most overlooked topic in PCB design is trace-to-trace separation for circuits of different voltages. Inadequate trace spacing may cause failures during ESD and transient testing, but also impact safety standards compliance and long-term reliability. Most designers are familiar with worldwide safety agency standards (UL, IEC, for example) which primarily address electrical hazards and flammability for the purpose of protecting the operator. There

are several insulation categories including operational, basic, supplemental, and reinforced, with each having an associated creepage (surface separation) distance dependent upon voltage, materials used, and the application environment. Circuits which share a common ground or are not separated by a hazardous potential fall into the operational insulation category. Operational insulation is the least restrictive category, and while it defines minimum creepage distances, it allows non-compliance if the offending traces are shorted together and no hazardous condition is created. The circuit can fail, but if no shock or fire hazard is created the spacing violation is accepted.

The point is these standards address safety issues, not reliability. It is quite easy to design a safety standards compliant product which lacks adequate trace separation for long term reliability. Depending on the operating environment, a voltage potential between traces can cause metallic whisker (dendrite) growth through the process of electromigration over a period of months or years. The rate of growth primarily depends on distance, voltage, humidity, and contaminants in the air or on the PCB surface. Eventually the dendrite will short the traces leading to product failure. PCB design standard IPC-2221, "Generic Standard on Printed Board Design", defines trace spacing requirements to avoid electromigration. IPC-2221 are generally less restrictive than operational insulation spacing, but there are no exceptions for violations. While it is possible to test around the operational insulation spacing requirements, the distance should not be reduced below the IPC-2221 guidelines.

Once the PCB design guidelines are understood it is time to consider circuit design flexibility and alternate assembly options so that corrective action can be taken without requiring a PCB design change.

Passive EMI Filtering

A preemptive design example can be shown for conducted EMI in the Ethernet interface. The impact of differential current to common mode current conversion, previously discussed, may be mitigated significantly by using a small CM inductor, commonly referred to as a balun. The CM inductor forces the currents in each line of the differential pair to be equal and opposite through transformer action (assuming a coupling coefficient of 1.00 and ignoring magnetizing current). The CM impedance is very high, and the differential impedance is (ideally) zero. The CM inductor can be placed in the layout and used if needed. If not needed, it may be omitted from the BOM and jumpers placed to short across the windings.

While the passive EMI solution provides a convenient alternative construction should the need arise it is not without some disadvantages. The CM choke is not ideal. It has low, but non-zero, differential inductance with affects both insertion loss and characteristic impedance. However, in most applications the impact is small and usually acceptable.

Active EMI Filtering

An alternative to the passive EMI filtering methods is an active solution for Ethernet applications. The Kinetic Technologies KTA1550 and KTA1552 active CM EMI and ESD suppressor ICs may be placed and routed on a PWB as a contingency. Upon later system

evaluation the parts and supporting components may be populated, or not, as needed with no further impact other than to the BOM. There is no need for jumpers or dummy components to complete the interconnect. The active solution does not suffer from the parasitic elements of the passive solution discussed above, plus it addresses ESD protection as a bonus. Figure 13 shows a typical KTA1550 through-routing of the data lines. The IC is located between the PHY and the LAN transformer.

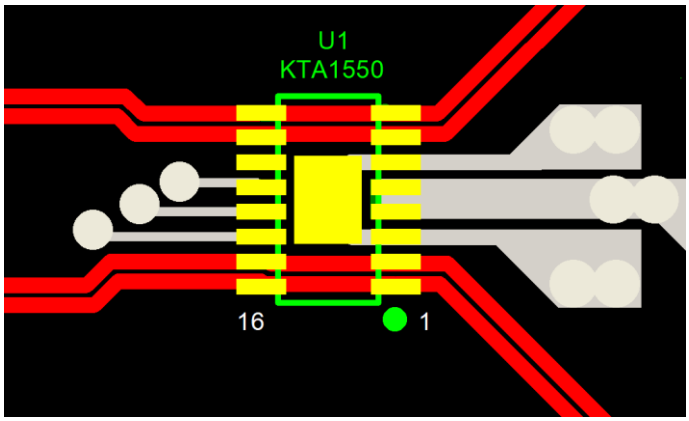

Figure 13. Transmission line configuration data lines are through-routed on pins 1-16, 2-15, 7-10, and 8-9 in this KTA1550 example (4mm x 2.5mm 16-pin TDFN).

The KTA1550 supports two twisted pair interfaces whereas the KTA1552 supports four. Each channel of a KTA1550/2 reduces common mode (CM) emissions by up to 10 dB from 1 to 125 MHz with virtually no impact to insertion, return loss, or characteristic impedance. The KTA155x ICs are compatible with both voltage and current mode PHYs. Figure 14 shows typical CM rejection improvement using the KTA155x in conjunction with two commonly used LAN transformer configurations.

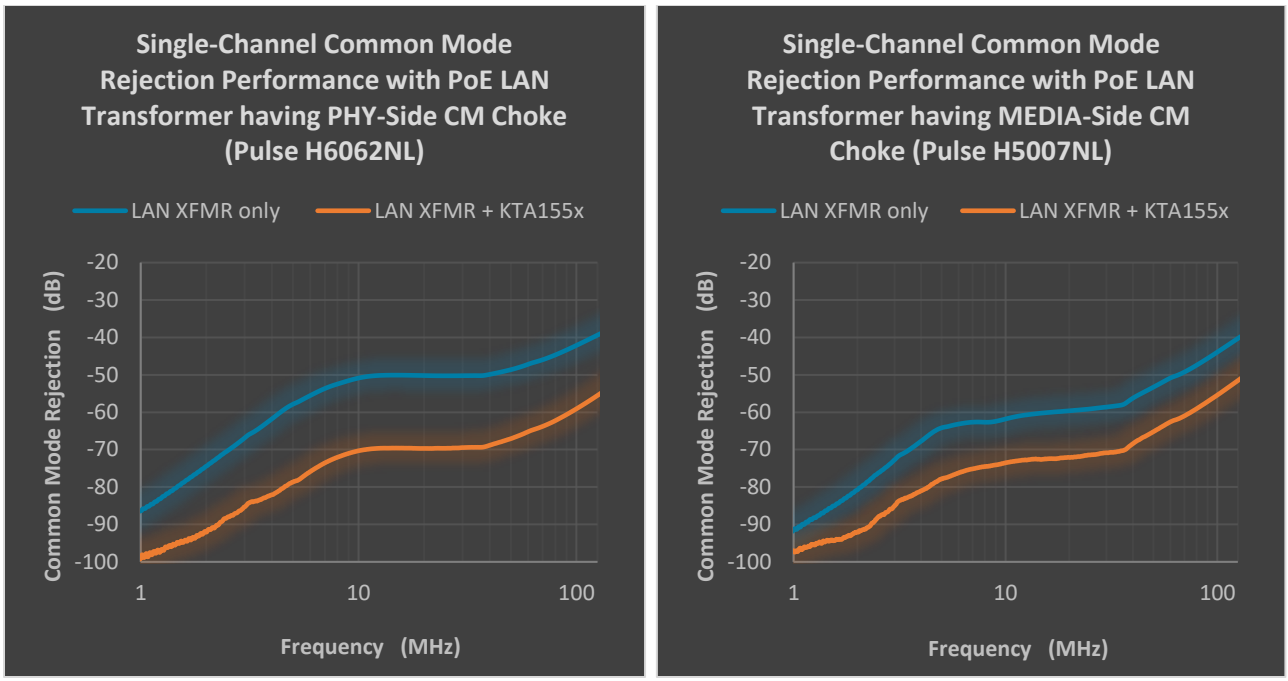

Figure 14. Typical Common Mode Rejection Improvement with KTA155x and PoE LAN Transformer having either MEDIA or PHY side CM chokes.

Figure 15 shows typical Insertion Loss Performance using the KTA155x in conjunction with either LAN transformer configuration. The test method measures difference in the PWB assembly with the KTA155x and support components removed and then again with those parts populated.

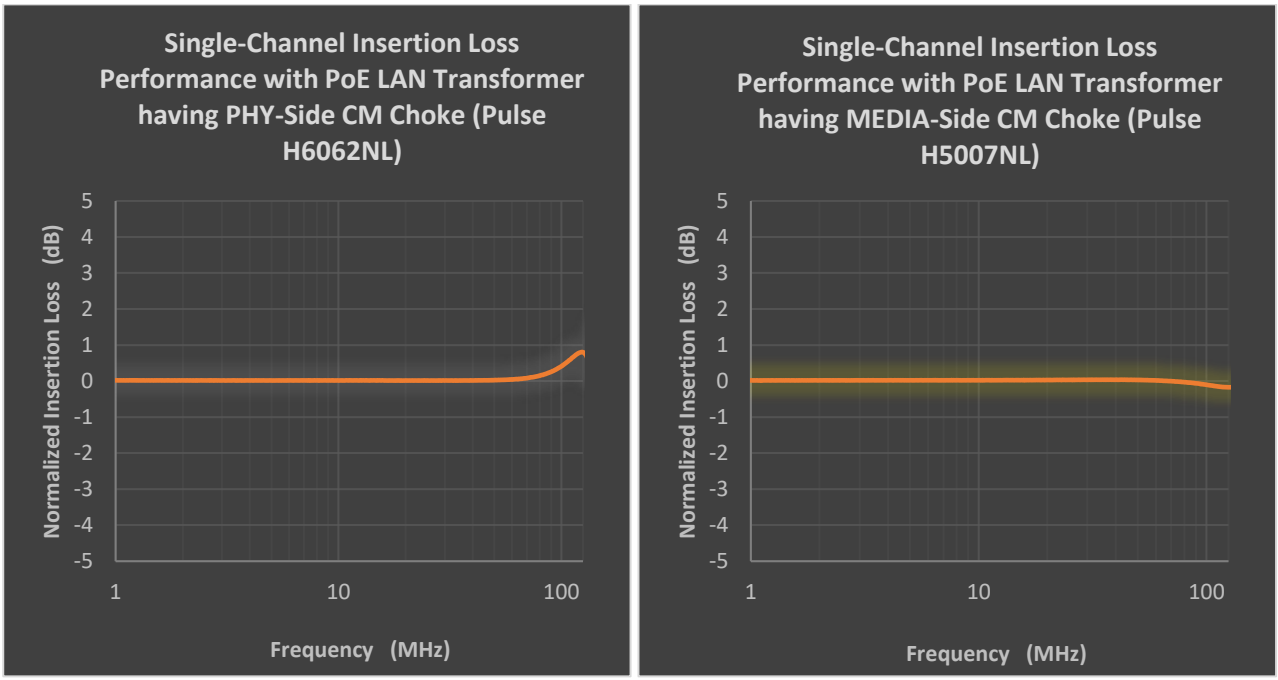

Figure 15. Typical Insertion Loss performance of KTA155x.

Figure 16 shows typical Return Loss Performance using the KTA155x in conjunction with either LAN transformer configuration.

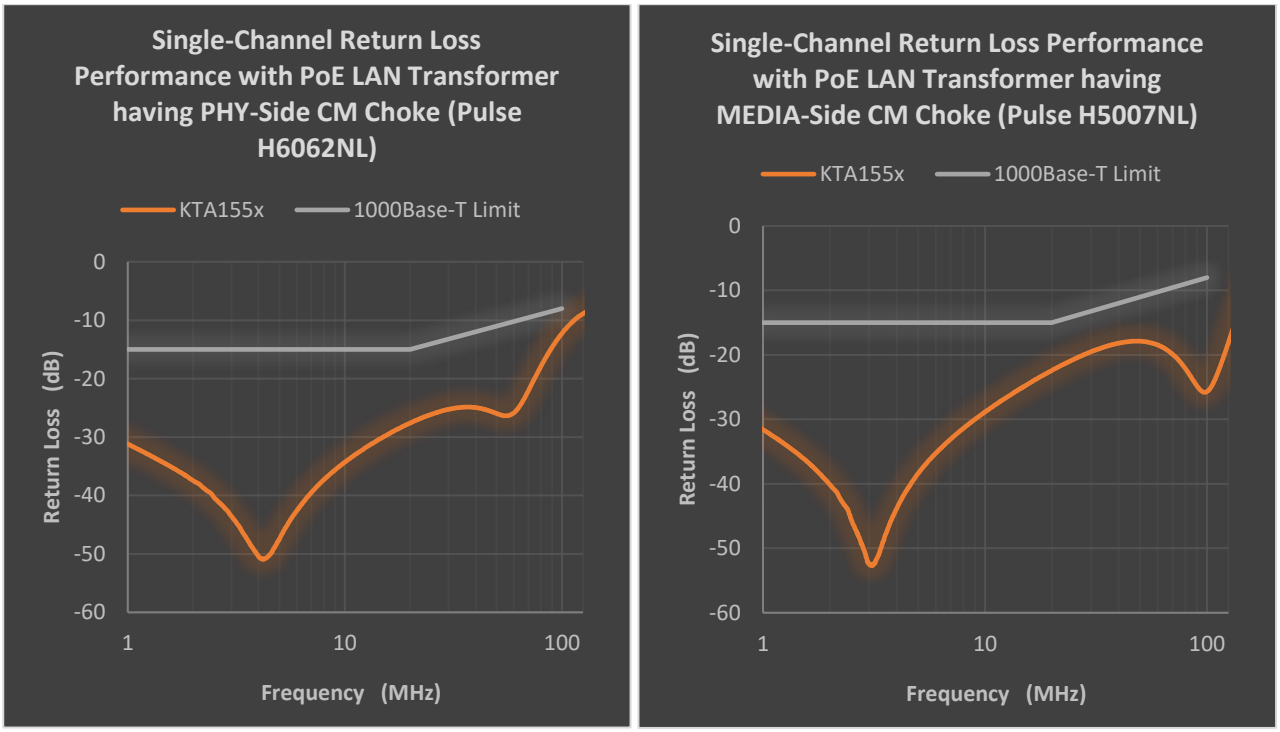

Figure 16. Typical Return Loss performance of KTA155x.

For more information on the test procedure used to generate these results please consult technical paper “KTA1550: How to Test AC Response of Active EMI & ESD Suppressor for Ethernet Applications”, Kinetic Technologies.

The results show that CM EMI emissions can be preemptively addressed just by including a layout option on the PCB for the KTA155x active CM EMI filter.

Conclusion

Many design issues may be preemptively addressed with proper prioritizing of signal routing, careful adherence to PCB design guidelines, and by including placement options for alternate circuitry for design flexibility. Some of these techniques have been discussed as they pertain to the specific example of Ethernet connected devices but are adaptable to other applications as well.

References

1. *High Speed Digital Design: Design of High Speed Interconnects and Signaling*, by Hangiao Zhang, Steven Krooswyk, and Jeff Ou, Elsevier Inc., 2015
2. *High Speed Digital Design: A Handbook of Black Magic*, by Howard Johnson and Martin Graham, Prentis-Hall, 1993
3. "KTA1550: How to Test AC Response of Active EMI & ESD Suppressor for Ethernet Applications", Kinetic Technologies Inc.
4. KTA1550 device datasheet, Kinetic Technologies Inc.
5. KTA1550-EVB, KTA1550 EVB schematic diagram and bill of materials, Kinetic Technologies Inc.
6. Application Note AN93: KTA1550 Design Guide, Kinetic Technologies Inc.