A HDMI design guide for successful high-speed PCB design

By Thomas Kugelstadt,
Senior Systems Engineer,
Texas Instruments Inc.

Introduction

This article presents design guidelines for helping users of HDMI mux-repeaters to maximise the device's full performance through careful printed circuit board (PCB) design. We'll explain important concepts of some main aspects of high-speed PCB design with recommendations. This discussion will cover layer stack, differential traces, controlled impedance transmission lines, discontinuities, routing guidelines, reference planes, vias and decoupling capacitors.

Layer stack

The pin-out of a HDMI mux-repeater is tailored for the design in HDTV receiver circuits (see Figure 1). Each side of the package provides a HDMI port, featuring four differential TMDS signal pairs, thus resulting in three input and one output port. The remaining signals comprise the supply rails, Vcc and ground, and lower speed signals such as the I2C interface, Hotplug-detect and the mux-selector pins.

A minimum of four layers are required to accomplish a low EMI PCB design (see Figure 2). Layer stacking should be in the following order (top-to-bottom): TMDS signal layer, ground plane, power plane and control signal layer.

- Routing the high-speed TMDS traces on the top layer avoids the use of vias (and the introduction of their inductances) and allows for clean interconnects from the HDMI connectors to the repeater inputs, and from the repeater output to the subsequent receiver circuit.
- Placing a solid ground plane next to the high-speed signal layer establishes controlled impedance for transmission line interconnects and provides an excellent low-inductance path for the return current flow.
- Placing the power plane next to the ground plane creates additional high-frequency bypass capacitance.
- Routing the slower speed control signals on the bottom layer allows for greater flexibility as these signal links usually have margin to tolerate discontinuities such as vias.

If an additional supply voltage plane or signal layer is needed, add a second power / ground plane system to the stack to keep it symmetrical. This makes the stack mechanically stable and prevents it from warping. Also the power and ground plane of each power system can be placed closer together, thus increasing the high-frequency bypass capacitance significantly.

Differential Traces

HDMI uses transition minimised differential signalling (TMDS) for transmitting high-speed serial data. Differential signalling offers significant benefits over single-ended signalling.

In single-ended systems, current flows from the source to the load through one conductor and returns via a ground plane or wire. The transversal electromagnetic wave (TEM), created by the current flow, can freely radiate to the outside environment causing severe electromagnetic interference (EMI) (Figure 3). Also noise from external sources induced into the conductor is unavoidably amplified by the receiver, thus compromising signal integrity.

Differential signalling instead uses two conductors, one for the forward, the other one for the return current to flow. Thus, when closely coupled, the currents in the two conductors are of equal amplitude but opposite polarity and their magnetic fields cancel. The TEM waves of the two conductors, now being robbed of their magnetic fields, cannot radiate into the environment. Only the far smaller fringing fields outside the conductor loop can radiate, thus yielding significantly lower EMI (see Figure 3).

Another benefit of close electric coupling is that external noise induced into both conductors equally appears as common-mode noise at the receiver input. Receivers with differential inputs are sensitive to signal differences only, but immune

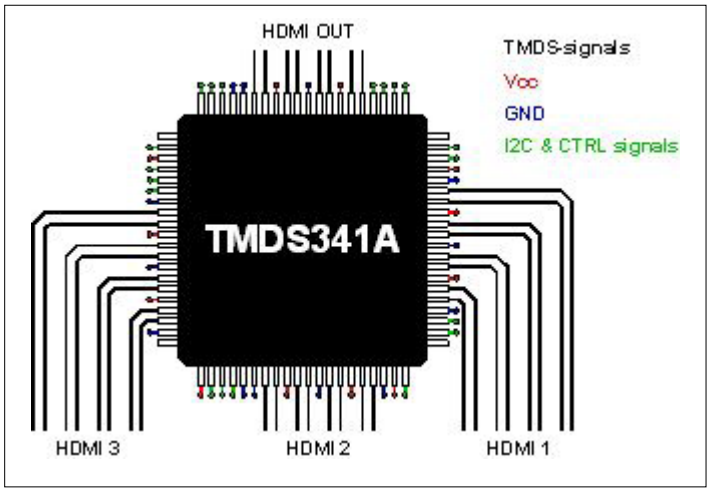

Figure 1: The device pin-out is tailored for HDTV receiver applications

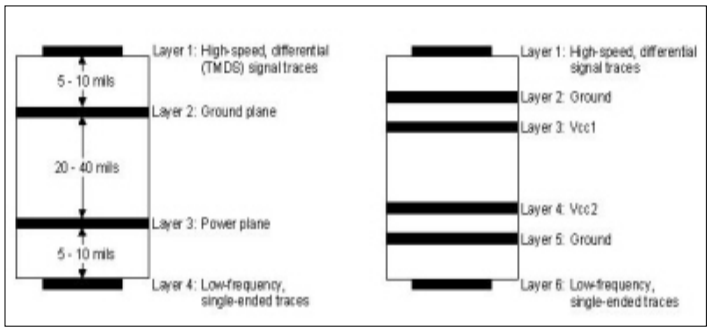

Figure 2: Recommended 4- or 6- layer stack for a receiver PCB design

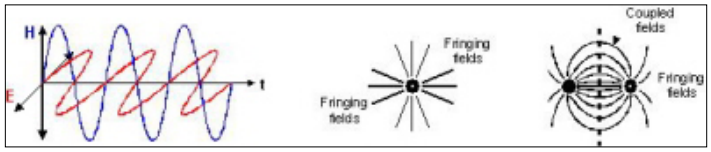

Figure 3: TEM wave radiation from the large fringing fields around a single conductor and the small fringing fields outside the closely coupled conductor loop of a differential signal pair.

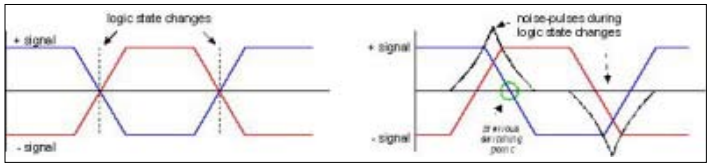

Figure 4: Traces of different electrical length cause phase shifts between signal, generating difference signals that cause serious EMI problems.

to common-mode signals. The receiver, therefore, rejects common-mode noise and signal integrity is maintained.

To make differential signalling work on a PCB, the spacing of the two traces of a differential signal pair must be kept the same across the entire length of the trace. Otherwise, variations in the spacing cause imbalances in the field coupling, thus, reducing the cancellation of the magnetic fields " leading to increased EMI.

In addition to larger EMI, changes in conductor spacing cause the differential impedance of the signal pair to change, thus creating discontinuities in an impedance-controlled transmission system, which leads to signal reflections compromising signal integrity.

Besides consistent spacing, both conductors must be of equal electrical length to ensure their signals reach the receiver inputs at the same time. Figure 4 shows the "+" and the """ signals of a differential pair during logic state changes for traces of equal and different length.

For traces of equal length both signals are equal and opposite. Therefore, their sum must add to zero. If the traces differ in electrical length, the signal on the shorter trace changes its state earlier than the one on the longer trace. During that time both traces drive currents into the same direction. Because the longer trace, which is supposed to act as return path, continues to drive current, the current of the "early" driving, shorter trace must find its return path via a reference plane (power or ground).

When adding both signals the sum signal diverts from the zero level during the transition phase. At high frequency these different signals appear as sharp transients of considerable magnitude, showing up on the ground plane, causing serious EMI problems.

Note that the width of the "noise" pulses is equal to the phase shift between the two signals, and can be translated into a time difference for a given frequency. This time difference, also known as intra-pair skew, is specified by HDMI for a receiver with 0.4 TBIT for a TMDS clock rate of 225 MHz, which translates to 178 ps maximum. For an HDMI transmitter the specification calls for 0.15 TBIT for a TMDS clock rate of 225 MHz, which translates to 66 ps maximum.

Because pixel generation requires the synchronous transmission of four differential TMDS signal pairs, (3 data + 1 clock), it must reach the receiver at the same time. Ideally, all four signal pairs should be of equal electrical length to ensure zero time difference. HDMI, however, allows for a maximum inter-pair skew, the time difference between signal pairs, for a receiver of 0.2 $T^{CHARACTER}$ + 1.78 ns, yielding a total of 2.67 ns for a TMDS clock of 225 MHz. For an HDMI transmitter, the specification calls for 0.2 $T^{CHARACTER}$ resulting in 888ps.

Controlled Impedance Transmission Lines

Controlled impedance traces are used to match the differential impedance of the transmission medium, e.g., cables, and the termination resistors. Differential impedance is determined by the physical geometries of the signal pair traces, their relation to the adjacent ground plane and the PCB dielectric. These geometries must be maintained across the entire trace length.

Figure 5 depicts the parameters relevant for impedance calculation for both, Microstrip traces (outer layer traces), and stripline traces (traces within the layer stack, typically sandwiched by two ground planes).

To calculate the trace geometries in Figure 5 for a 100 differential impedance TMDS signal pair, the closed-form equations 1 " 6 can be applied.

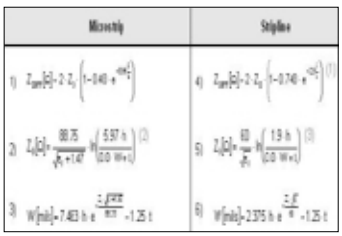

Figure 5: Physical geometries of differential traces

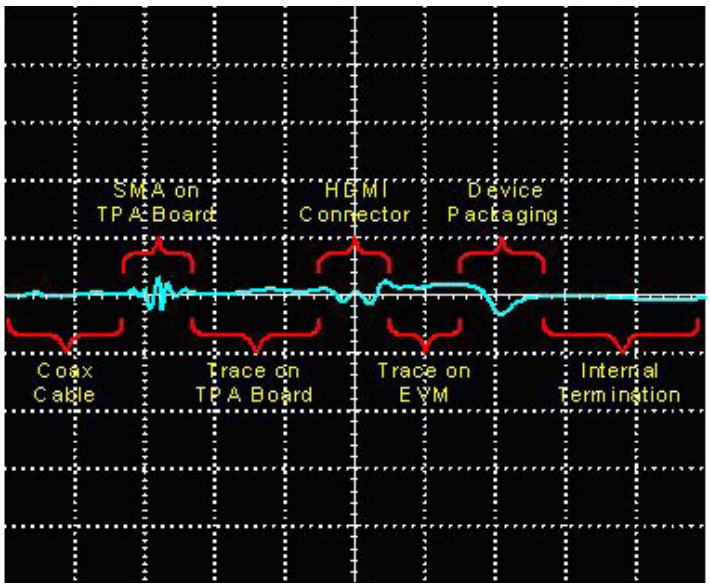

Figure 6: TDR display revealing the locations of discontinuities

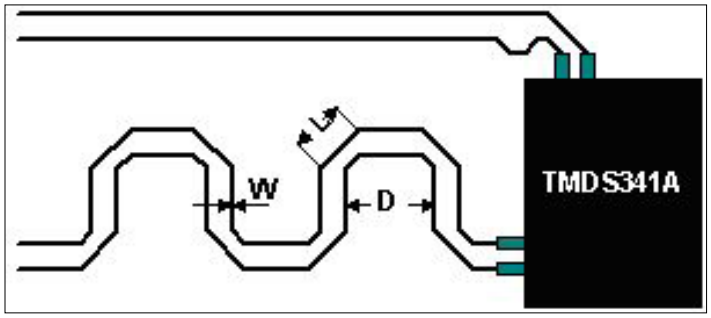

Figure 7: Skew reduction via meandering using chamfered corners

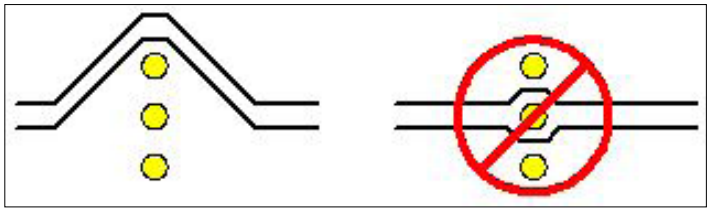

Figure 8:Routing around an object

1. For loosely coupled striplines, s > 12 mils, the number 0.748 might be replaced with 0.374
2. For W < 2h the maximum error is 3%
3. For best accuracy keep b " t > 2W and b > 4t, where b is the dielectric thickness between ground planes

With regards to the distance between a different signal pair and its environment, Figure 5 shows a trace X that is not associated with the current flow in the adjacent "+" and """ conductors. X can be a trace of another signal pair, a ground shield trace or a TTL/CMOS trace.

For adjacent signal pairs and shield traces, make distance d 3 s. Running the shield trace (preferably ground) on one side potentially creates an imbalance that can increase EMI. Ground trace shields should have a scattering of vias to the underlying ground plane.

Note: At first glance the equations above represent an inexpensive way to attain the

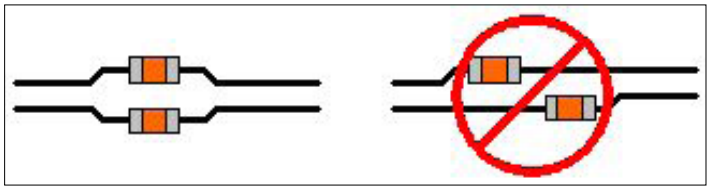

Figure 9: Lumping discontinuities

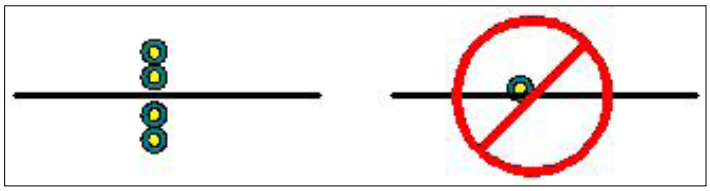

Figure 10: Avoiding via clearance sections

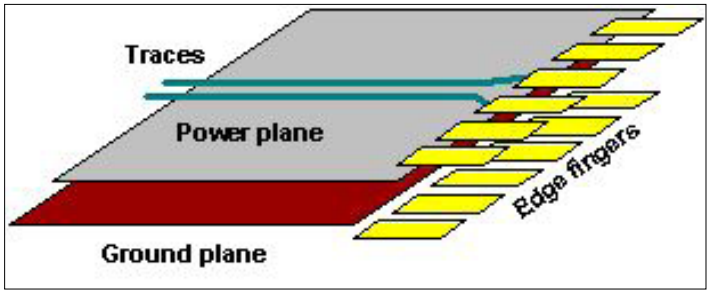

Figure 11: Keeping planes out of the area between edge-fingers

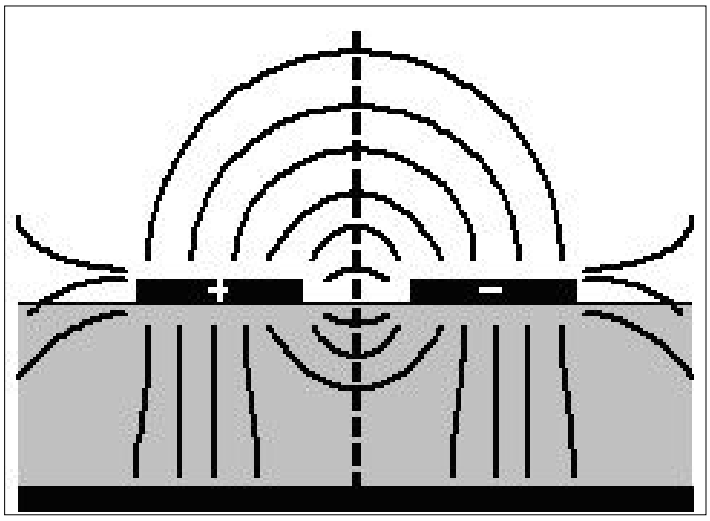

Figure 12: Field coupling within a microstrip structure

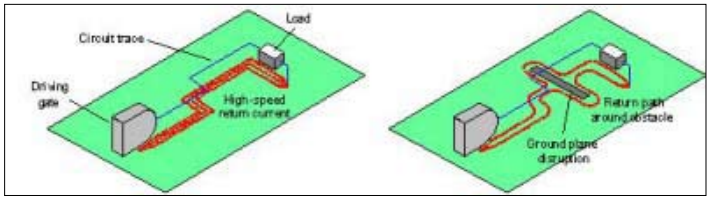

Figure 13: Return current paths in solid versus slotted ground planes

trace geometries. However, these functions are based on empirical data and represent good approximations at best. The actual accuracy might vary significantly and various sources even cite possible errors of up to 10 per cent.

A more accurate, and in the long term cheaper approach, is to use a 2D or better field solver. This is a software tool that solves Maxwell's Equations and calculates the electric and magnetic fields for an arbitrary cross-section transmission line. From these, it also calculates the electrical performance terms, such as characteristic impedance, signal speed, crosstalk and differential impedance. Some field solvers can also calculate the current distributions inside conductors. The advantage a 2D field solver wields over an approximation is the flexibility to consider almost any arbitrary cross-section geometry. In addition to the first-order terms such as line width, dielectric thickness and dielectric constant, second-order terms such as trace thickness, solder mask and trace etch back can be considered.

Discontinuities

Discontinuities are locations in the signal path where the differential trace impedance deviates from its specified value (of 100 15 per cent for HDMI), and assumes either higher or lower impedance values. Discontinuities cause signal reflections due to impedance mismatch compromising signal integrity. These are primarily the result of changes in the effective trace width or in the line-to-line spacing caused either by unavoidable transitions in the trace geometries along the signal path, or by poor routing of the signal traces.

Potential locations for discontinuities are:

- where the solder pads of the HDMI connector meet the signal traces
- where signal traces meet vias, component pads of resistors, or IC-pins
- 90o bends in signal traces
- where a signal pair is split to route around an object

Discontinuities are detected during differential impedance, TDR, tests. A TDR, (time-domain reflectometer), is an electronic instrument used to characterise and locate faults in metallic conductors.

A TDR transmits a fast rise time pulse along the conductor. If the conductor is of uniform impedance and properly terminated, the entire transmitted pulse will be absorbed in the far-end termination and no signal will be reflected back to the TDR. But where impedance discontinuities exist, each discontinuity will create an echo that is reflected back to the reflectometer (hence the name). Increases in the impedance create an echo that reinforces the original pulse while decreases in the impedance create an echo that opposes the original pulse.

The resulting reflected pulse that is measured at the output/input to the TDR is displayed or plotted as a function of time and, because the speed of signal propagation is relatively constant for a given transmission medium, can be read as a function of trace length.

The goal in PCB design must be to minimise discontinuities wherever possible, thus eliminating reflections and maintaining signal integrity. Following a minimum set of routing guidelines helps avoiding unnecessary discontinuities. The remaining, unavoidable discontinuities should be lumped, that is their areas should be kept small and placed together as close as possible. The idea is to concentrate the points of reflection to a certain area rather than having them distributed across the entire signal path.

The magnitude of the discontinuities seen using a TDR are directly effected by the edge rate of the pulse used by the TDR. The faster the TDR edge the more discontinuities will show up, and the larger the impedance spike will appear. With the HDMI specification they have defined the edge rate that is to be used to be 200ps. Figure 6 illustrates this point. The lower line on the graph was taken using a 30ps edge rate and the upper line was taken with a 200pf filter. The discontinuities created by the SMA launch onto the TPA board that show up on the low line are completely invisible when the 200ps edge rate filter is applied.

Routing Guidelines

Guidelines for routing PCB traces are necessary when trying to maintain signal integrity and lower EMI. Although there seems to be an endless number of pre-

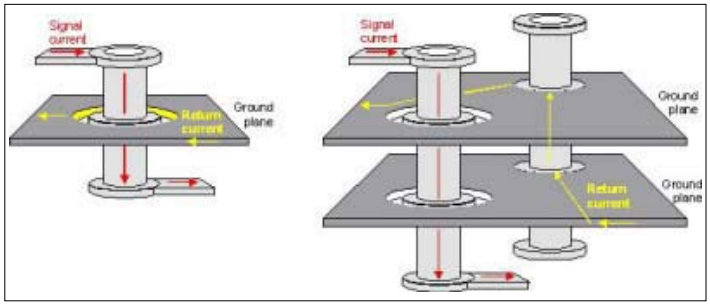

Figure 14: Return current paths for a single and a multiple layer change

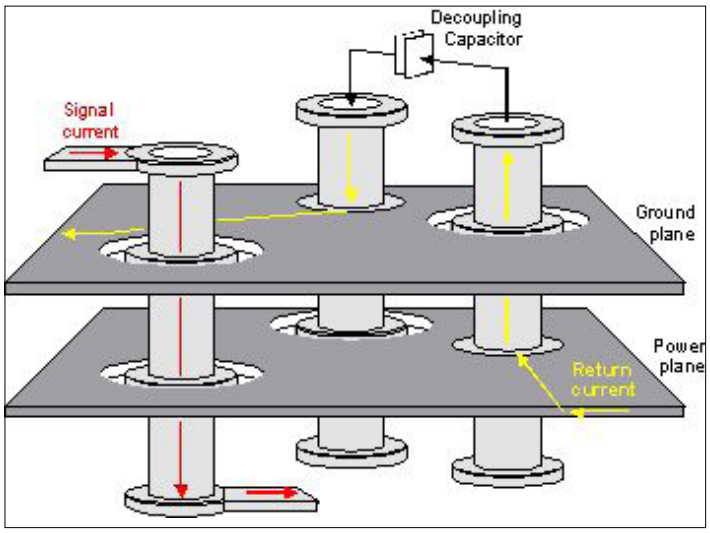

Figure 15: Return current paths for a single- and a multiple layer change

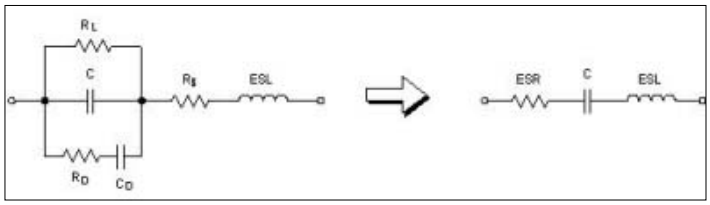

Figure 16: Capacitor losses modelled by a series resonance circuit

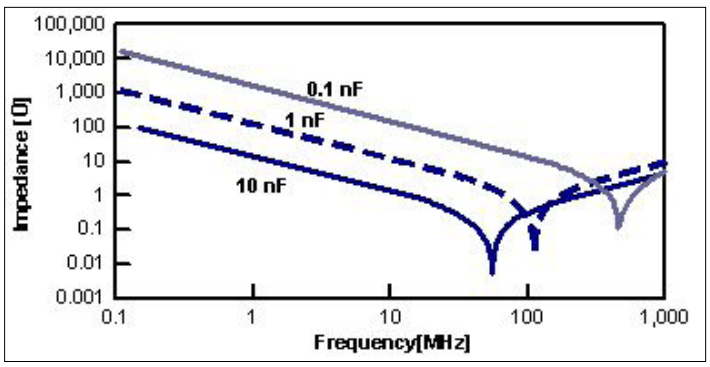

Figure 17: Capacitor impedance versus frequency

cautions to be taken, this section provides only a few main recommendations as layout guidance.

1. Reduce intra-pair skew in a differential trace by introducing small meandering corrections at the point of mismatch.
2. Reduce inter-pair skew, caused by component placement and IC pinouts, by making larger meandering correction along the signal path. Use chamfered corners with a length-to-trace width ratio of 3 to 5. The distance between bends should be at least 8 to 10 times the trace width.
3. Use 45o bends (chamfered corners), instead of right-angle (90o) bends. Right-angle bends increase the effective trace width, which changes the differential trace impedance creating a small discontinuity. A 45o bends is seen as an even smaller discontinuity.
4. When routing around an object, route both trace of a pair in parallel. Splitting the traces changes the line-to-line spacing, thus causing the differential impedance to change and discontinuities to occur.
5. Place passive components within the signal path, such as source-matching resistors or ac-coupling capacitors, next to each other. Routing as in case a) does create wider trace spacing than in b); however, the resulting discontinuity is limited to a far narrower electrical length.
6. When routing traces next to a via or between an array of vias, make sure that the via clearance section does not interrupt the path of the return current on the ground plane below.
7. Avoid metal layers and traces underneath or between the pads of the HDMI connectors for better impedance matching. Otherwise they may cause the differential impedance to drop below 75 and fail your board during TDR testing.
8. Use the smallest size possible for signal trace vias and HDMI connector pads as they have less impact on the 100 differential impedance. Large vias and pads can cause the impedance to drop below 85.
9. Use solid power and ground planes for 100 impedance control and minimum power noise.
10. For 100 differential impedance use the smallest trace spacing possible, which usually is specified by your PCB vendor. Make sure that the geometries in **Figure 5** are: s < h, s < W, W < 2h, and d > 2s. Even better, use a 2D field solver to determine the trace geometries more accurately.
11. Keep the trace electrical length between the HDMI connector and the device as short as possible to minimise attenuation.
12. Use good HDMI connectors whose impedances meet the specifications.
13. Place bulk capacitors, (i.e., 10 ¼F), close to power sources, such as voltage regulators or where the power is supplied to the PCB.
14. Place smaller 0.1 ¼F or 0.01 ¼F capacitors at the device.

Reference Planes

The power and ground planes of a high-speed PCB design usually must satisfy a variety of requirements. At DC and low frequencies they must deliver stable reference voltages, such as Vcc and ground, to the supply terminals of integrated circuits and termination resistors.

At high frequencies reference planes, and in particular ground planes, serve numerous purposes. For the design of controlled impedance transmission systems, the ground plane should provide strong electric coupling with the differential traces of an adjacent signal layer. As mentioned earlier, close coupling causes the magnetic fields to cancel, thus minimising EMI through reduced TEM wave radiation of the remaining fringing fields. To accomplish close coupling, place the ground plane next to a high-speed signal layer.

Although differential signalling ideally does not require separate current return paths, there always will be some form of common-mode noise currents capacitively coupling into the closest reference plane, (which ideally is a ground plane).

Providing a continuous, low-impedance return path for these current requires the reference planes to be of solid copper sheets, free from voids and crevices.

Layer stacks with multiple power systems can benefit from reference planes that are stitched with vias. Here ground planes of different layers are connected

through a number of vias placed at regular intervals across the board. Similar is valid for the stitching of like power planes.

For stitched reference planes, it is important that the via clearance section (or anti pads in the case of ground vias) does not interfere with the path of the return current. In the case of an obstacle the return current will find its way around it. However, by doing so, the current's electromagnetic fields will most likely interfere with the fields of other signal traces introducing crosstalk. Moreover, this obstacle will adversely affect the impedance of the traces passing over it.

Vias

The term via commonly refers to a plated hole in a printed circuit board. While some applications require through-hole vias to be wide enough to accommodate the leads of through-hole components, high-speed board designs mainly use them as trace routing vias when changing signal layers, or as connecting vias to connect SMT components to the required reference plane and also to connect reference planes of the same potential to each other (recall the via stitched ground plane of the previous section).

Layers connecting to a via do so by making direct contact with a pad surrounding the via (the via pad). Layers that must not connect are separated from the via by a clearance ring. Every via has capacitance to ground, which can be approximated using the following equation:

$$C = \frac{1.41 \cdot \varepsilon_r \cdot T \cdot D_1}{D_2 - D_1}$$

where D2 = diameter of clearance hole in ground planes, [in].
D1 = diameter of pad surround via, [in].
T = thickness of printed circuit board, [in].
ε_r = dielectric constant of the circuit board.
C = parasitic via capacitance, [pF].

Because the capacitance increases proportional with size, trace vias in high-speed designs should be as small as possible to avoid signal degradation caused by heavy capacitive loading.

When connecting decoupling capacitors to a ground plane, or interconnecting ground planes, the via inductance becomes more important than its capacitance. The magnitude of this inductance is approximately:

$$L = 5.08 \cdot h \cdot \left[\ln\left(\frac{4 \cdot h}{d}\right) + 1 \right]$$

where L = via inductance, [nH].
h = via length, [in].
d = via diameter, [in].

Because this equation involves a logarithm, changing the via diameter does little to influence the inductance. A big change may be effected by changing the via length or by using multiple vias in parallel. Therefore, connect decoupling capacitors to ground by using two paralleled vias per device terminal. For low inductance connections between ground planes, use multiple vias in regular intervals across the board.

Although it is highly recommended not to change layers of high-speed traces, if the necessity still occurs ensures a continuous return current path. Figure 14 on the left shows the flow of the return current for a single layer change and on the right for a multiple layer change.

The ability for the current flow to change from the bottom to the top of the ground plane is provided by a metallic laminate of the inner clearance ring. Thus, when a signal passes through a via and continues on the opposite side of the same plane, a return current discontinuity does not exist.

Changing a signal trace from one layer to another by crossing multiple reference planes complicates the design of the return current path. In the case of two ground planes, a ground-to-ground via must be placed near the signal via to ensure a continuous return current path, (right diagram in Figure 14). If the reference planes are of different voltage potentials, such as the power and ground planes in Figure 15, the design of the return path becomes messy as it requires a third via and a decoupling capacitor. The return current flow begins at the bottom of the power plane, where it is closest to the signal current. It then flows through the power via, across the decoupling capacitor into the ground via, and returns on top of the ground plane.

Current return paths comprising multiple vias and decoupling capacitors possess high inductance, thus compromising signal integrity and increasing EMI. If possible, avoid changing layers during high-speed trace routing, as it usually worsens board performance, complicates design and increases manufacturing cost.

Decoupling Capacitors

Decoupling capacitors provide a local source of charge for ICs requiring a significant amount of supply current in response to internal switching. Insufficient decoupling causes a lack of supply current required, which may prevent the IC from working properly resulting in signal integrity data errors to occur. This requires them to provide low impedance across the frequency range of interest. To accomplish that, a common approach is to distribute an array of decoupling capacitors evenly across the board. In addition to maintaining signal integrity, decoupling capacitors serve as EMC filters preventing high-frequency RF signals from propagating throughout the PCB.

When connecting a capacitor between the power and ground planes we are actually loading the power supply with a series resonant circuit, whose frequency dependent R-L-C components represent the equivalent circuit of a real capacitor. Figure 16 shows the parasitic components of an initial equivalent circuit and their conversion into a series resonant circuit.

The leakage resistance RL represents the loss through leakage current at low frequencies. RD and CD indicate the losses due to molecular polarisation, (RD), and dielectric absorption, (CD). RS depicts the resistance in the leads and the plates of the capacitor. The three resistive losses are combined into one equivalent series resistance (ESR). As in the ESR case, the equivalent series inductance (ESL) combines the inductance of the capacitor plates and the internal leads.

Note that the capacitor connecting vias, although low in impedance, contribute a significant amount to the series inductance. Therefore, reduce via inductance by using two vias per capacitor terminal.

Figure 17 shows the progression of capacitor impedance (Z) versus frequency for a 10 nF capacitor. At frequencies far below the self resonance frequency (SRF), the capacitive reactance is dominant. Closer to SRF the inductive reactance gains influence trying to neutralise the capacitive component. At SRF the capacitive and inductive reactance cancel and only the ESR is effective. Note that the ESR is frequency dependent, and in contrast to popular believe, does not reach its minimum at SRF. The impedance Z however does.

The reason why the paralleling of capacitors in a distributed decoupling network works is because the total capacitance increases to , where n is the number of decoupling capacitors used. And with , the capacitor impedance is reduced to for frequencies below SRF. Similar holds true for the inductance. Here and because the impedance decreases to for frequencies above SRF.

Designing a solid decoupling network must include lower frequencies down to DC, which requires the implementation of large capacitors. Therefore, to provide sufficient low impedance at low frequencies, place 1 ¼F to 10 ¼F tantalums at the output of voltage regulators and at the point where power is supplied to the PCB. For the higher frequency range place several 0.1 ¼F or 0.01 ¼F ceramics next to every high speed switching IC.

Summary
Without claiming to be complete, the objective of this paper is to cover the main aspects of high-speed PCB design. Despite the enormous amount of technical literature, seminars, newsletters and internet forums on the subject, this document's intent is to provide PCB designers with design guidelines in a very comprehensive way.

Following the recommendations presented will help accomplishing an EMC-compliant board design in the shortest time possible.