Physical Implementation of the High-Speed Design Process

David Wiens
Mentor Graphics Corporation

Introduction

Layout designers play a critical role in the design of high-speed circuits. In an optimized process, they take constraints defined by engineering and create a layout that adheres to those design rules. Typically, however, the process is more ad hoc with constraints being transferred via paper or voice, putting more responsibility on the layout designer to translate the requirements and implement a functional design. In some organizations, engineering has taken over layout of high-speed signals to ensure accurate implementation of constraints. Unfortunately, this often results in designs that are impossible to complete or are un-manufacturable.

Today's complex designs can have advanced constraints on 80-90% of the board's signals, significantly increasing the layout difficulty and design cycle time. To appropriately manage and implement these constraints, layout designers must understand concepts such as differential impedance/signaling, star net topologies, interconnect delay and crosstalk coupling. This paper will address:

- High-speed constraint categories in "electrical" and "physical" incarnations
- An optimized design process focusing on collaboration between engineering and layout departments
- Routing best practices to adhere to advanced high-speed constraints
- The impact/benefit of advanced fabrication technologies like HDI and embedded passives on high-speed design.

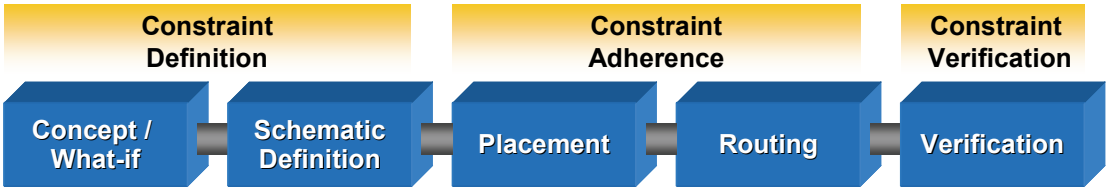

Figure 1 - Typical Design process, From Concept Through Final Verification.

High-speed process overview

Understanding the high-speed design process boils down to understanding the constraints used to define performance requirements, as well as techniques used to ensure adherence and verification of those constraints. Constraints enable efficient transfer of design intent from engineers to layout designers, eliminating the need for "voice-driven mouse" (i.e. engineer-over-shoulder) or paper napkin drawings. A typical design process is shown in Figure 1. It is overlaid with the stages constraints go through during the design process, from definition to adherence to final verification. If constraints are defined early in the process – during conceptualization and schematic creation – they can be utilized throughout the rest of the process. Tools later in the process (like placement and routing) must adhere to the constraints to minimize the number of design iterations required if problems are found during the final design verification. This discussion focuses on an understanding of the constraints typically used, and how they apply to the layout stage of the design process.

Constraint definition and adherence

Impedance constraints

Constraints for single line or differential impedance can be specified by engineering in either 'electrical' or 'physical' terms. Electrical constraints are specified in Ohms, while physical constraints are specified as trace width/layer combinations. Definition of a materials stack-up is critical to the accurate definition, adherence and verification of impedance constraints. Most design tools provide field solvers to calculate impedance automatically based on the stack-up. The charts in Figure 2 illustrate the relationship of the layer stack-up to trace impedance.

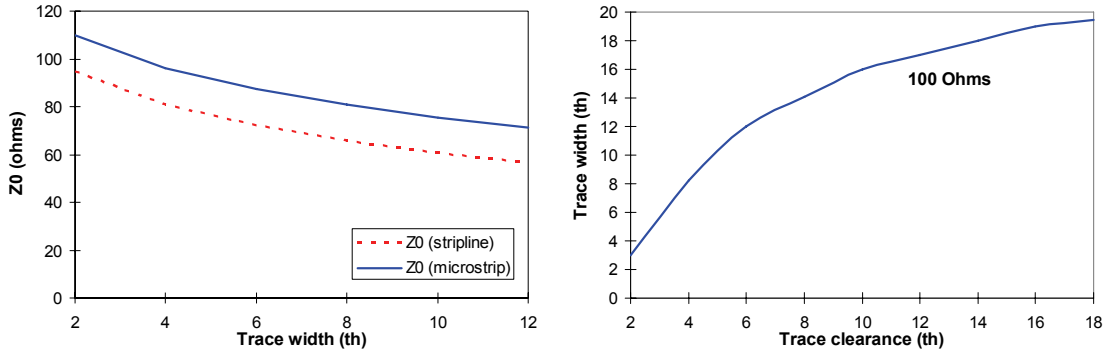

(2a) *(2b)*

Figure 2 - Impact of trace width on signal impedance. Figure 2a shows how impedance drops as trace width increases. The change is different for a stripline configuration (trace on a layer sandwiched between two plane layers) and a microstrip configuration (trace on outer layer with plane layer on just one side). Figure 2b shows how trace width and clearance align to maintain constant differential impedance.

Timing constraints

Constraints for timing also have equivalent electrical (i.e. delay) and physical (i.e. length) formats. Engineers define timing constraints for individual nets, or for matched sets of nets to minimize timing skew. Net topology is also critical to effective timing management. A PCB layout with advanced topologies in use today is shown in Figure 3. These topologies have created the need for "virtual pins", which float between actual component pins to enable constraint definition for properly timed nets. Figure 4 highlights the difference in signal quality between a traditional chained topology and an improved star topology.

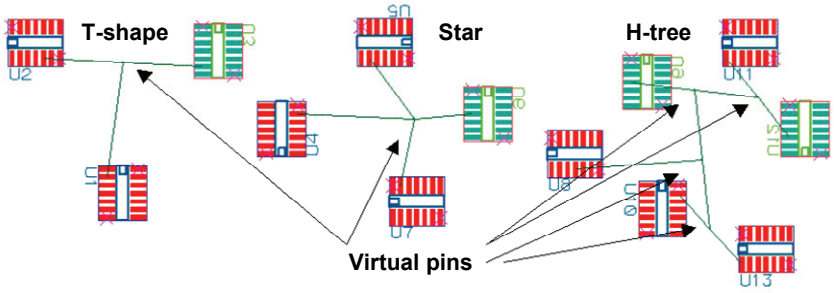

Figure 3. Advanced net topologies used to define signal timing.

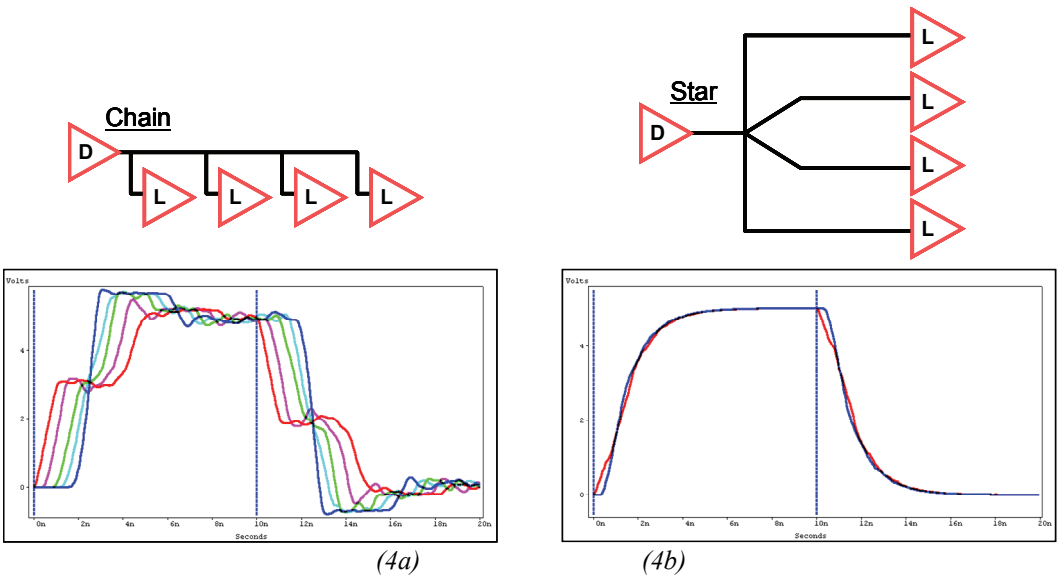

(4a) *(4b)*

Figure 4. The impact of net topology on signal quality. Figure 4a shows a chained topology and resulting waveform – note the timing skew and degraded quality due to reflections. Figure 4b shows a star topology where all loads received the signal at the same time, resulting in a much cleaner waveform.

Net tuning is the process of routing a net (or set of nets) to meet the timing constraints. If a net is too short, extra length (typically in the form of serpentine patterns) must be added. Designers employ multiple types of tuning patterns (see Figure 5), depending on space availability and concern about inductive self-coupling within serpentine patterns.

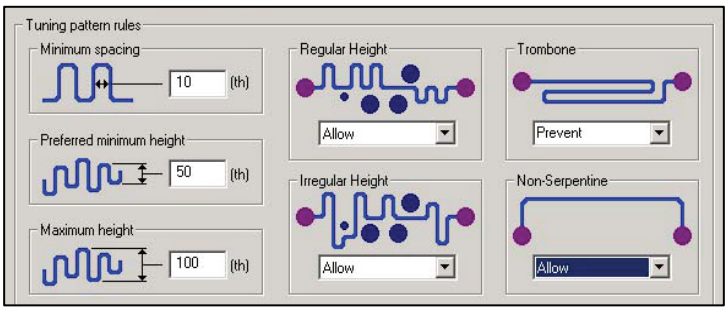

Figure 5: Constraints for various tuning structures

Differential pair constraints

Differential signals have increased in popularity largely due to multi-gigabit data transmission technologies that have replaced traditional parallel bus structures with a single (differential) serial structure. Constraints for these topologies include: pair tolerance (how much timing skew is allowed between the traces – can be as little as 0.025% maximum allowed length difference); convergence and separation lengths (the lengths at which ideal and worst-case coupling is allowed); and the spacing between the differential traces. Timing constraints on differential pairs are also typical. Figure 6 shows three differential signals whose routing was very closely coupled, with a timing requirement that resulted in added length in the form of serpentine patterns.

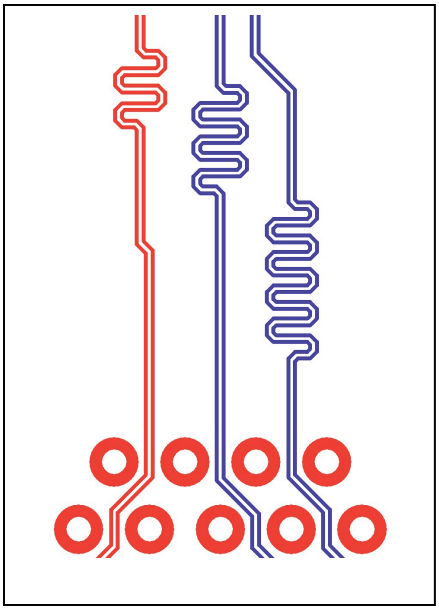

Figure 6 - Tightly coupled routing for tuned differential signals

Crosstalk constraints

Constraints for crosstalk can be specified in either electrical terms (maximum millivolts of coupling) or physical terms (typically a parallelism table of allowed coupling lengths at various trace separations). Electrical terms are much easier to define, but require additional information (e.g. driver edge rate and voltage swing).

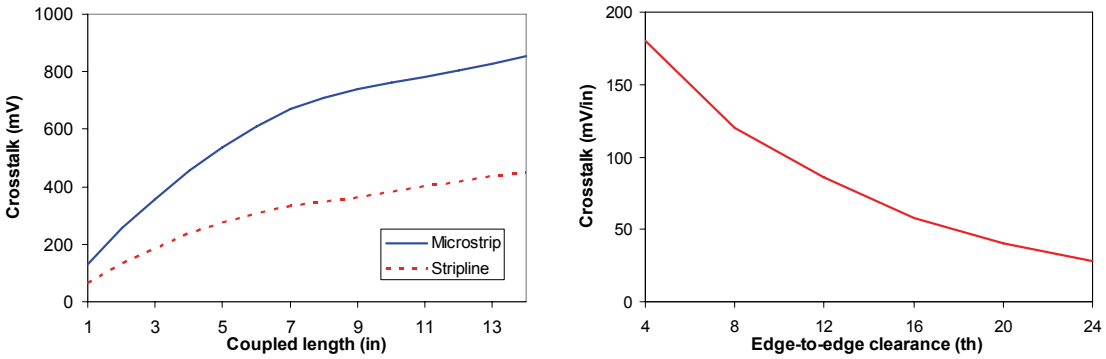

(7a) *(7b)*

Figure 7 - The effects of parallelism and trace clearance on crosstalk. Figure 7a shows the increase in crosstalk relative to coupled length (parallelism) – note that traces in stripline configurations have less crosstalk than microstrips. Figure 7b shows decreased crosstalk with increased trace-to-trace separation.

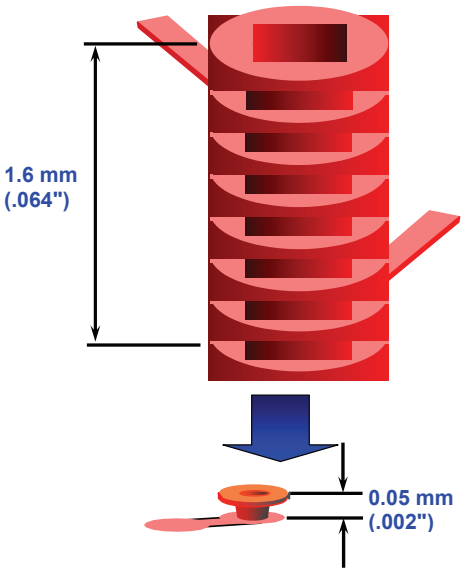

Figure 8: Comparison of through-hole via to HDI via

Routing best-practices
This section will highlight typical issues found routing high-speed signals, and propose some options to evaluate.

Tune critical nets last
The typical approach to routing critical signals is to route and tune them first, then fix/lock them in place so they can't be modified. However, these signals can then significantly block the routes of other signals. If you have a high volume (>70%) of constrained nets, route them first, but don't tune or lock them until last. This ensures that they have room to connect, but they can still be pushed around to provide space for other, non-critical nets.

Always try to get electrical constraints
Electrical constraints will always be more accurate than physical constraints. This is because physical constraints are an interpretation of electrical constraints, and are thus subject to conversion error. For instance, if a timing constraint is specified as length the actual delay may vary widely depending on which layer the signal is routed on (e.g. a 2.5in constraint could yield a 356ps delay on layer 1, or a 422ps delay on layer 6 – almost a 20% difference!). Be aware that electrical constraints require additional data to model the layer stack-up and component I/O, but creation of this information can greatly improve the quality of placement and routing to high-speed constraints.

Be aware that external layers are typically faster
External (microstrip or embedded microstrip) layers are generally faster than internal (stripline) layers. This can come in handy when routing matched nets where the Manhattan (un-routed) lengths are not equal. If the longest net is routed on an external layer, the resulting delay that must be matched will be shorter than on internal layers. This can obviously improve the performance of the matched set, but will also result in a significant savings of routed copper. Note that most designers keep critical routes off the external layers because they are not shielded, and thus subject to EMI. However, this principle can still be utilized whenever there is a difference in the velocity of propagation between any two layers.

Advanced PCB fabrication and high-speed design
Some relatively new advances in fabrication technology designed to enable product miniaturization can also have a positive impact on the electrical performance of a circuit.

High-density interconnect (HDI)
Vias generated with HDI (microvia) technology can have significantly smaller parasitics (capacitance and inductance) and delay compared to through-hole vias because of smaller holes and fewer via pads (see Figure 8). They don't create via stubs which can act as antenna at high frequencies (e.g. create a through-hole via for a ten layer board, then just use it to route from layer one to layer three – the unused portion of the via can act as an antenna). Circuit EMI is also reduced because HDI has thinner layers, allowing the signals to be closer to the ground planes, and because distributed capacitance is available.

Embedded passives

Embedded passives (resistors, capacitors and inductors) can also yield performance improvements. The passives can be placed on internal layers, directly underneath the associated active device, minimizing inductance and reducing delay. In addition, embedded resistor technology allows creation of nearly any resistance, improving line impedance matching capability to minimize reflections. Embedded passives also reduce crosstalk, EMI and noise – switching noise can be greatly minimized due to the proximity of decoupling capacitors to active devices.

Conclusion

High-speed constraints are consuming today's designs. To optimize product performance and team productivity, layout designers must increase their collaboration with design engineers. To do this, they should have a firm grasp on high-speed constraint types, as well as techniques to ensure adherence during placement and routing. The alternative is an increase in the number of design cycles due to circuit rejections during final verification or physical prototype test. Advanced fabrication technologies also play a role in high-speed design, enabling significant improvements in system performance. Layout designers who understand high-speed and advanced fabrication technologies can ensure that the PCB design meets and exceeds the product's performance specifications.