10 Fundamental Rules of High-speed PCB Design, Part 2

Beyond Design

by Barry Olney, IN-CIRCUIT DESIGN PTY LTD / AUSTRALIA

In last month's column, I introduced the 10 fundamental rules of high-speed PCB design (Figure 1). The first rule was to establish design constraints before commencing the design. This prime strategy sets constraints upfront based on pre-layout analyses or recommendations and guidelines and is integral to the design flow to maintain the established requirements. This month, I will elaborate on the importance of controlling the impedance and floor planning the placement based on connectivity.

II. Control the Impedance: Match the transmission line impedance to the driver and load. Create the stackup and define terminations to match the impedance.

For perfect energy transfer, the impedance of the driver must match the transmission line—assuming there is a high-impedance load. A good transmission line is one that has constant impedance along the entire length of the line, so no mismatches result in reflections.

Digital design typically uses a characteristic impedance of 50–60 ohms. However, this value becomes more critical as the edge rates increase. Different technologies also have specific impedance requirements. For example, Ethernet is 100 ohms, USB is 90 ohms differential, DDR2 memory is 50/100, and DDR3/4 is 40/80 single-ended/differential impedance. Thus, controlling impedance simultaneously on each signal layer with many different technologies can become a challenge. Further, as operating voltages decrease, the associated

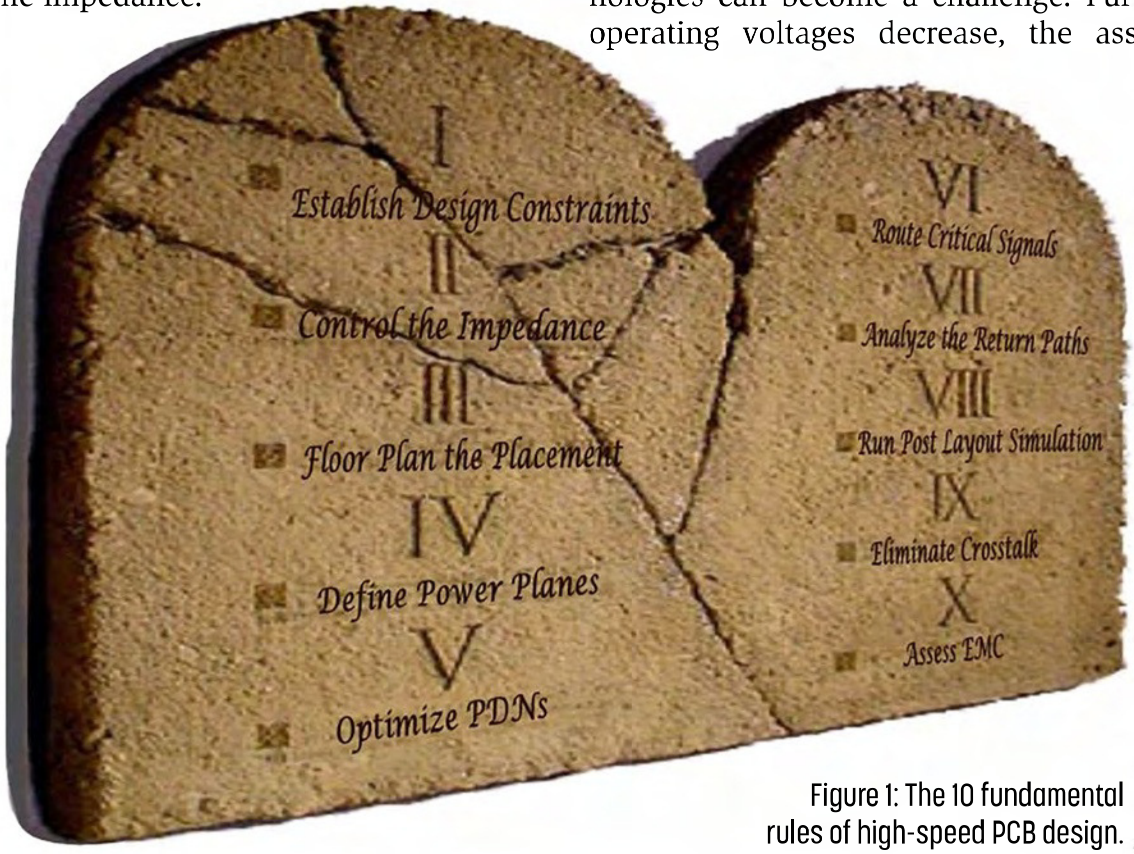

Figure 1: The 10 fundamental rules of high-speed PCB design.

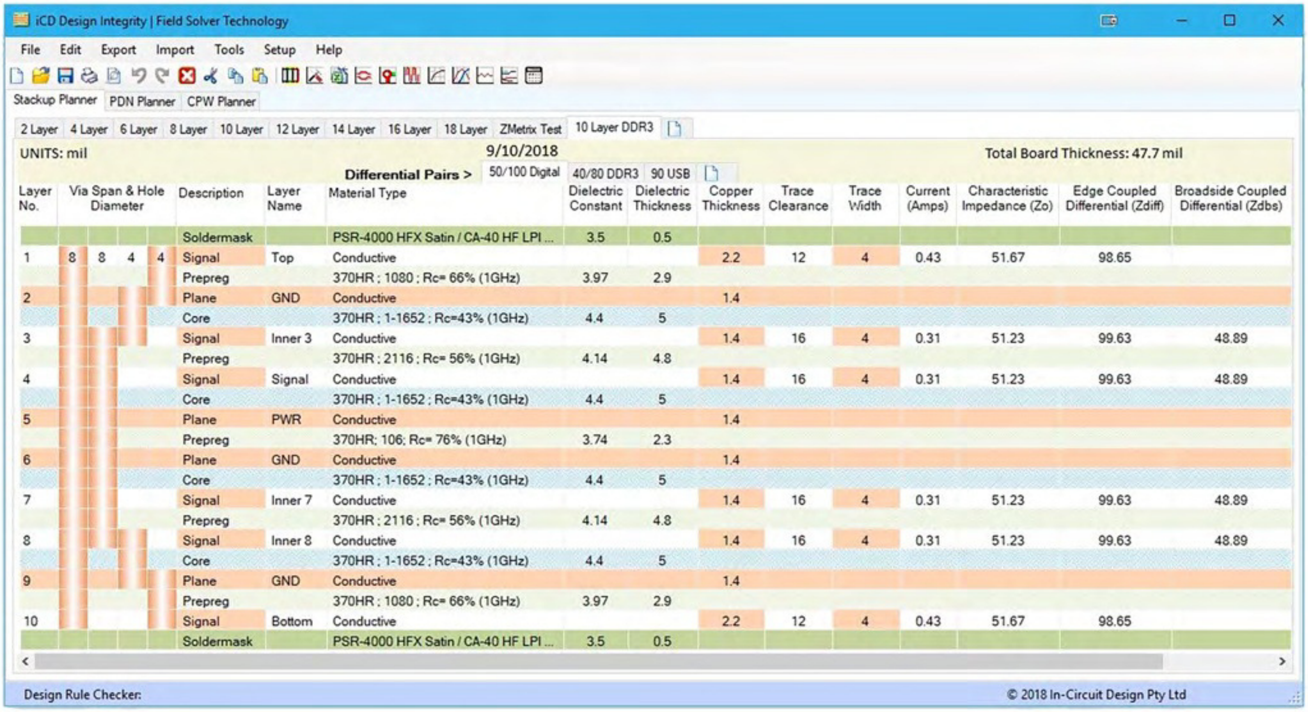

UNITS: mil

9/10/2018

Total Board Thickness: 47.7 mil

Layer No.	Via Span & Hole Diameter	Description	Layer Name	Material Type	Dielectric Constant	Dielectric Thickness	Copper Thickness	Trace Clearance	Trace Width	Current (Amps)	Characteristic Impedance (Zo)	Edge Coupled Differential (Zdiff)	Broadside Coupled Differential (Zdbs)
		Soldermask		PSR-4000 HFX Satin / CA-40 HF LPI ...	3.5	0.5							
1	8 8 4 4	Signal	Top	Conductive			2.2	12	4	0.43	51.67	98.65	
		Prepreg		370HR ; 1080 ; Rc= 66% (1GHz)	3.97	2.9							
2		Plane	GND	Conductive			1.4						
		Core		370HR ; 1-1652 ; Rc=43% (1GHz)	4.4	5							
3		Signal	Inner 3	Conductive			1.4	16	4	0.31	51.23	99.63	48.89
		Prepreg		370HR ; 2116 ; Rc= 56% (1GHz)	4.14	4.8							
4		Signal	Signal	Conductive			1.4	16	4	0.31	51.23	99.63	48.89
		Core		370HR ; 1-1652 ; Rc=43% (1GHz)	4.4	5							
5		Plane	PWR	Conductive			1.4						
		Prepreg		370HR; 106; Rc= 76% (1GHz)	3.74	2.3							
6		Plane	GND	Conductive			1.4						
		Core		370HR ; 1-1652 ; Rc=43% (1GHz)	4.4	5							
7		Signal	Inner 7	Conductive			1.4	16	4	0.31	51.23	99.63	48.89
		Prepreg		370HR ; 2116 ; Rc= 56% (1GHz)	4.14	4.8							
8		Signal	Inner 8	Conductive			1.4	16	4	0.31	51.23	99.63	48.89
		Core		370HR ; 1-1652 ; Rc=43% (1GHz)	4.4	5							
9		Plane	GND	Conductive			1.4						
		Prepreg		370HR ; 1080 ; Rc= 66% (1GHz)	3.97	2.9							
10		Signal	Bottom	Conductive			2.2	12	4	0.43	51.67	98.65	
		Soldermask		PSR-4000 HFX Satin / CA-40 HF LPI ...	3.5	0.5							

Figure 2: 10 Layer stackup with matched Ethernet, DDR3, and USB impedances (iCD Stackup Planner).

noise margins also decrease, which makes it even more critical to match the impedance. Figure 2 shows differential pairs set up to accommodate three different technologies on the same layers of the substrate.

Notice how the signal traces are tightly coupled to the reference planes. This helps prevent unwanted radiation, particularly on the outer microstrip signals. The center dielectric material (between layers 5 and 6) is also very thin (2.3 mils) and provides low-impedance planar capacitance to the power distribution networks (PDNs).

Unfortunately, drivers do not have the same impedance as the transmission line (typically 10–35 ohms), so terminations are used to balance the impedance, match the line, and minimize reflections. Reflections occur whenever the impedance of the transmission line changes along its length. This can be caused by unmatched drivers/loads, layer transitions, different dielectric materials, stubs, vias, connectors, and integrated-circuit (IC) packages. By understanding the causes of these reflections and eliminating the source of the mismatch, a design can be engineered with reliable performance.

Figure 3 shows how using a 12-mA LVCMOS 1.8-V driver of a Spartan 6 FPGA and an 18.7-ohm series resistor is required to match the driver to the 51.67-ohm trace on the outer

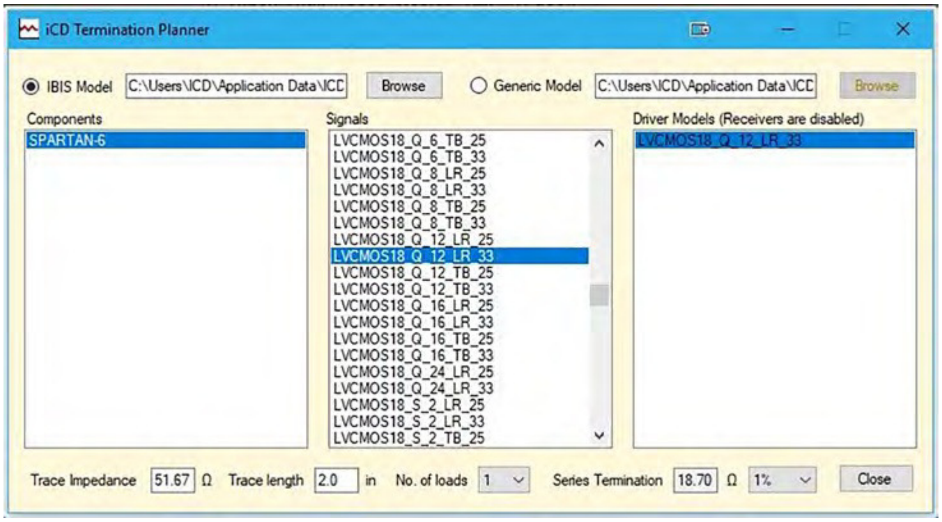

● IBIS Model C:\Users\ICD\Application Data\ICD

○ Generic Model C:\Users\ICD\Application Data\ICD

Components
SPARTAN-6

Signals
LVCMOS18_Q_6_TB_25
LVCMOS18_Q_6_TB_33
LVCMOS18_Q_8_LR_25
LVCMOS18_Q_8_LR_33
LVCMOS18_Q_8_TB_25
LVCMOS18_Q_8_TB_33
LVCMOS18_Q_12_LR_25
LVCMOS18_Q_12_LR_33
LVCMOS18_Q_12_TB_25
LVCMOS18_Q_12_TB_33
LVCMOS18_Q_16_LR_25
LVCMOS18_Q_16_LR_33
LVCMOS18_Q_16_TB_25
LVCMOS18_Q_16_TB_33
LVCMOS18_Q_24_LR_25
LVCMOS18_Q_24_LR_33
LVCMOS18_S_2_LR_25
LVCMOS18_S_2_LR_33
LVCMOS18_S_2_TB_25

Driver Models (Receivers are disabled)
LVCMOS18_Q_12_LR_33

Trace Impedance 51.67 Ω Trace length 2.0 in No. of loads 1 Series Termination 18.70 Ω 1%

Figure 3: Matching the Spartan 6 driver to the transmission line (iCD Termination Planner).

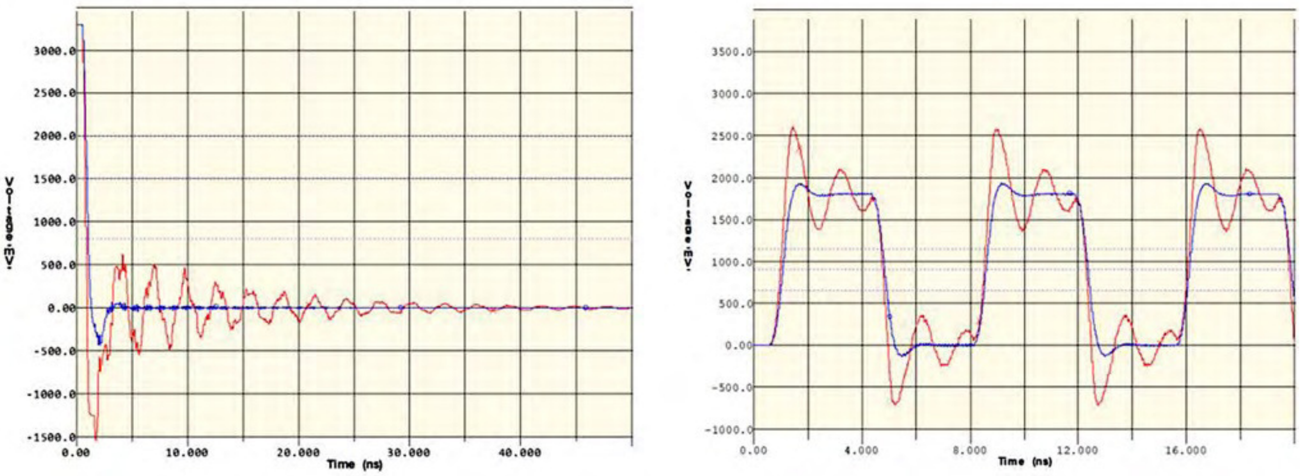

Figure 4: Ringing is reduced dramatically by adding a series terminator (simulated in HyperLynx).

layer. This is automatically derived from the IV curves of the Spartan 6 IBIS model by the iCD Termination Planner.

Figure 4 illustrates the ringing (red) in an unmatched transmission line. This ringing, which is also represented by over/undershoot (right), is dramatically reduced by terminating the transmission line with an 18.7-ohm series resistor (blue). Controlling the impedance of the transmission lines ensures that your product will perform more reliably and exhibit improved signal quality and reduced crosstalk and electromagnetic radiation.

III. Floor Plan the Placement Based on Connectivity: Place components by functionality and analog and digital groups to minimize interaction between different logic families and improve routability and timing.

Since aggressor signals induce crosstalk onto the victim signal, the higher the aggressor voltage, the more crosstalk will be induced. Therefore, it is best to segregate groups of nets according to their signal amplitude. This strategy prevents higher voltage nets (e.g., 3.3 V) from affecting lower voltage nets (e.g., 1.5 V), which have lower noise margins.

It is also preferable to partition these groups by rise time and frequency. Position the fastest devices closest to the connector (Figure 5). The placement should be graduated in descending order of speed down to the analog sections farthest from the connector to avoid noise coupling into sensitive devices.

All analog signals should be routed in the analog section and all the digital signals in the digital section. Of course, control signals must route between them. The segregation method, which I have used for many years, is to employ route fences or keep-outs. Route fences can be defined by placing elongated keep-outs on all layers. They are placed to direct the routing. No signal can cross these fences on any layer.

It is essential to keep in mind that high-speed return currents follow the path of least inductance. For example, if a trace is routed from the digital-to-analog converter (DAC) to

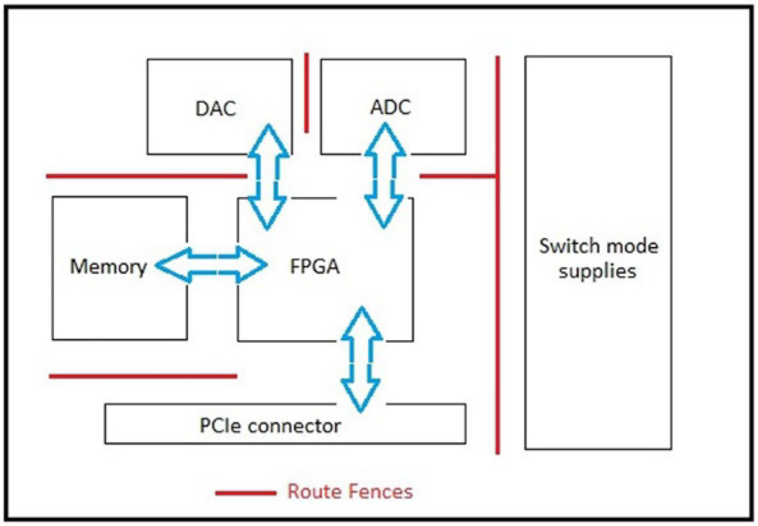

Figure 5: Components are segregated by logic groups and functionality.

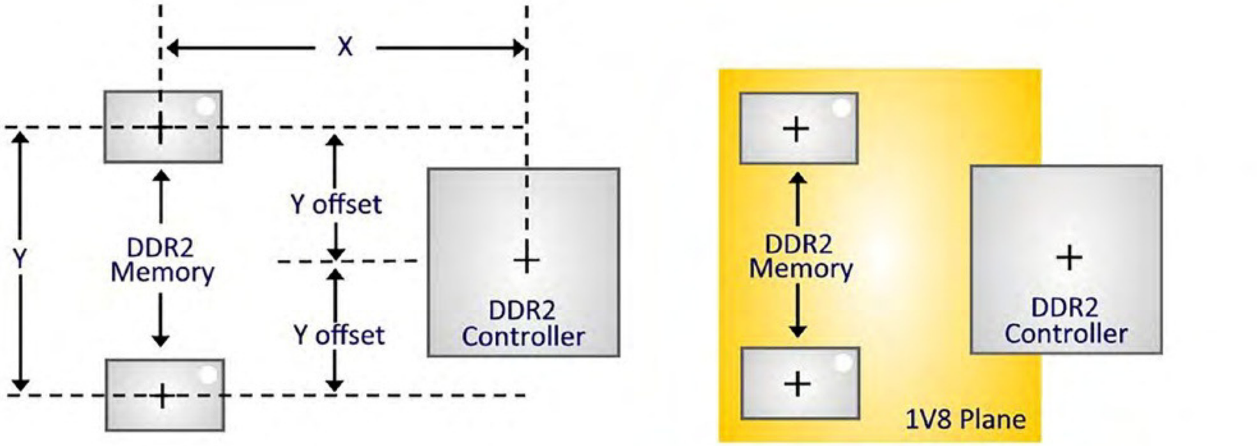

Figure 6: Processor and memory device placement requirements and 1.8-V plane.

the field-programmable gate array (FPGA), the return current path will be directly below that trace and will not wander into nearby sections. Route fences also control the autorouter by preventing signals from crossing and allowing the control signals to pass. Additionally, a disconcerting problem with high-speed boards is that their failure mode may be intermittent behavior across multiple manufacturing runs. In this case, the proper layout of the PCB may mean the difference between a reliable product and a board that performs intermittently.

Moreover, flight time delay and skew are key concerns in high-speed PCB design. One of the driving factors for flight time and skew performance is the placement of components. Controlling the maximum placement of devices, combined with the assumption that good design practices are adhered to, limits the maximum signal delay to approximately the longest Manhattan (X + Y) distance of the signals contained in a specific clock domain. Why the longest Manhattan distance? This is due to skew matching requirements. All of the shorter nets in a clock domain must be lengthened to skew match to the longest run length. Therefore, flight time and skew for an entire clock domain are governed by the maximum placement, along with the routing rules that constrain the matching of the trace lengths.

In the classic high-speed design flow, timing specifications and simulation results are compared to determine placement and routing constraints. Given a length constraint, a designer can manage signal integrity by controlling the PCB trace topology of the various parts of an interface.

Figure 6 (left) shows the required placement for T-topology routing of a DDR2 controller and memory chips. The purpose of the placement guide is to limit the maximum trace lengths and allow for routing and via space, which can be a challenge. This placement does not restrict whether these devices are placed on the top or bottom of the board. The region of the board used for DDR2 circuitry must be isolated from other signals. The DDR2 keep-out region is defined for this purpose and shown in Figure 6 (right). The 1.8-V power plane should cover this entire region, and non-DDR2 signals should be kept out of this region. Controlling the placement of devices minimizes interaction between different logic families, limits maximum trace length, reduces flight time delay and skew, and assists in complying with timing specifications.

Key Points:

- The impedance of the driver must match the transmission line for perfect energy transfer
- Digital design typically uses a characteristic impedance of 50–60 ohms; however, different technologies have specific impedance requirements
- Associated noise margins decrease as operating voltages decrease

- Drivers do not have the same impedance as the transmission line, so terminations are used to match the impedance
- Reflections occur whenever the impedance of the transmission line changes along its length
- Ringing can be dramatically reduced by terminating the transmission line
- The higher the aggressor voltage, the more crosstalk will be induced into the victim signal
- Position the fastest devices closest to the connector and be graduated in descending order of speed down to the analog sections farthest from the connector
- Route fences direct the routing, which prevents signals from crossing and allows the control signals to pass
- The return current path will be directly below that trace and will not wander into nearby sections
- The proper layout of the PCB may mean the difference between a reliable product and a board that performs intermittently
- Controlling the maximum placement of devices limits the maximum signal delay to approximately the longest Manhattan distance
- Given a length constraint, a designer can manage signal integrity by controlling the PCB trace topology of the various parts of an interface

Further Reading

1. *Beyond Design: Stackup Planning, Part 2* by Barry Olney, *The PCB Design Magazine*, July 2015.

2. *Beyond Design: Controlled Impedance Design* by Barry Olney, *The PCB Design Magazine*, May 2015.

3. *Beyond Design: Impedance Matching–Terminations* by Barry Olney, *The PCB Design Magazine*, October 2013.

4. *Beyond Design: Critical Placement* by Barry Olney, *The PCB Magazine*, September 2012.

5. *Beyond Design: Mixed Digital-Analog Technologies* by Barry Olney, *The PCB Magazine*, August 2012.

6. *High-Speed Digital Design: A Handbook of Black Magic, First Edition* by Howard Johnson and Martin Graham, Prentice Hall, 1993.

Barry Olney is managing director of In-Circuit Design Pty Ltd (iCD), Australia, a PCB design service bureau that specializes in board-level simulation. The company developed the iCD Design Integrity software incorporating the iCD Stackup, PDN, and CPW Planner. The software can be downloaded from www.icd.com.au. To read past columns or contact Olney, click here.

New NIST Method Measures 3D Polymer Processing Precisely

Recipes for 3D printing, or additive manufacturing, of parts have required as much guesswork as science. Until now.

Now, researchers at the National Institute of Standards and Technology (NIST) have demonstrated a novel light-based atomic force microscopy (AFM) technique—sample-coupled-resonance photorheology (SCRPR)—that measures how and where a material's properties change in real time at the smallest scales during the curing process.

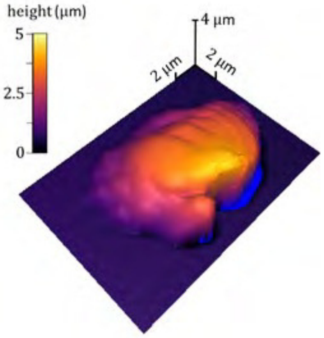

3D printing, or additive manufacturing, is lauded for flexible, efficient production of complex parts but has the disadvantage of introducing microscopic variations in a material's properties. Because software renders the parts as thin layers and then reconstructs them in 3D before printing, the physical material's bulk properties no longer match those of the printed parts. Instead, the performance of fabricated parts depends on printing conditions.

NIST's new method measures how materials evolve with submicrometer spatial resolution and submillisecond time resolution—thousands of times smaller-scale and faster than bulk measurement techniques.

Surprising the researchers, interest in the NIST technique has extended well beyond the initial 3D printing applications. Companies in the coatings, optics and additive manufacturing fields have reached out, and some are pursing formal collaborations, NIST researchers say.

(Source: NIST)