

Chris Mangelsdorf

## VDD is NOT your BFF:

### *Four things you didn't know could go wrong*


Somewhere on your schematic, there is a small circle. It's probably labeled VDD or AVDD or something equally innocuous. But that harmless-looking circle embodies the biggest lie in the industry (see Figure 1). Nobody knows the full extent of the damage that circle causes, but a conservative guess is that it's responsible for at least 90% of silicon failures and redesigns. That circle is your power supply. And it is almost certainly doing you wrong.

If you think about it, representing a complex multidimensional power delivery network (PDN) with an ideal voltage source is the wildest “approximation” we make. In an age of detailed transistor modeling and powerful multifaceted tools, where every simulation is precise beyond our wildest dreams, assuming that a power supply is perfect is a whopper of an approximation. Yet we do it every day without thinking.

Sure, you *know* about power supply noise and PSRR. You dutifully check your design with an ac source at some point. But there are four things you don't know, and those four things can work together to create a perfect storm of interference on your supply rails.

#### Threat #1: Bandini Mountain

If you're a transistor sitting on the silicon, looking out at the world, you see a dizzying array of interconnections and parasitics that are your



**FIGURE 1:** A real power supply. The VDD in your simulation doesn't look like this, does it? Ever wonder if maybe your simulations aren't telling you the whole story?

power supply, your lifeline to the circuit board for the volts and amps that keep you alive. You'll notice, though, that your power supply has two general features. First, the current you're receiving must pass through a complex series of metal traces, vias, bond wires, bumps, laminate traces, etc. that are all basically inductive. Second, there is quite a bit of capacitance in your neighborhood. Most notably there are some decoupling capacitors hanging around to protect you from all the inductance in your lifeline and to swallow all the noise of your neighbors. It's a happy and contented scene, and you're grateful to be a transistor in such a thoughtful well-designed chip.

You—as a contented transistor—are the victim of a dangerous illusion. The inductors you're looking at are basically grounded (in an ac sense) on the board. The capacitors around you also have a pretty good ground connection. So, from where you're sitting, the

VDD line is basically an inductor and a capacitor in parallel! [See Figure 2(a). Also, see “Resonance Refresher” if you are a little rusty when it comes to inductors.] That combination can ring, of course, and that's bad. But the scarier thing is that parallel resonance can have really high impedance. (“Really high” for a power supply is 1  $\Omega$ . Remember that SoCs can draw many amps.) That means the power supply is going to be choked off at some frequency! Yikes. The resonant spike probably won't cut off the supply altogether, but it is *the* dominant feature of the power supply impedance profile. So, your supply *will* be very noisy at the resonant frequency.

This nightmare is not only possible; it's unavoidable. It is so common that its impedance peak has a name: *Bandini Mountain* [2].

Now you might think that this name comes from the pioneering technologist who discovered this feature and warned the rest of us. Not so. It comes

from a publicity stunt by Bandini Fertilizer company in anticipation of the 1984 Olympics in Los Angeles. They created a 100-ft-high pile of manure and staged mock Olympic events on it for televised commercials. Thus, the name *Bandini Mountain* associates the dominant PDN impedance peak with a huge mound of doo-doo, hinting at the affection that packaging engineers have for it.

The most remarkable thing about VDD's Bandini Mountain is its predictability. While it may show up anywhere from 10 Mz to 100 MHz, in high performance packaging it's probably going to be in the 100–200 MHz range. Mysteriously, it doesn't seem to evolve over time with advances in technology and packaging. It sticks to 100 MHz like a curse. So, if you're going to have problems with your power supply, it's likely to happen around 100 MHz, give or take a factor of two.

Just to be clear, the Bandini phenomenon is primarily the work of two unavoidable IC features: package inductance and on-chip (mean-

ing "on-die") decoupling capacitance. (Note that ground will also have inductance. But it is typically less than what VDD has and can be included in the VDD total inductance without a loss of generality.) There will be several other minor hills caused by board inductances and other decoupling capacitors, but they occur at lower frequencies and can be engineered to be less of a nuisance. For the sake of simplicity, it is usually assumed that any board, laminate, or intentional package capacitance is doing its job, so the package inductance is terminated in an ac ground off-chip.

Fun fact: What looks like a high-impedance parallel resonance from the silicon viewpoint looks like a low-impedance series resonance from the board. Measuring this impedance dip from outside is the easiest way to measure the Bandini effect directly. Indirect methods are provided automatically by Murphy's Law and include such wonderful things as a noise hump in the skirts of your PLL spectrum at 100 MHz offset.

## Threat #2: Loop BADwidth

Having a resonance in your power supply is bad enough, but having it occur in the neighborhood of 100 MHz is awful. Why? Because feedback circuits don't have the bandwidth to deal with it.

Most of the analog circuits we care about use feedback. It's the best way to make accurate, robust signals, references, and biases. The self-correcting nature of feedback is also the very best defense against power supply noise or disturbances of any kind (unless, of course, the disturbance is coupled to the input). But as frequencies increase the loop gain begins to fall, which causes the feedback defenses to weaken.

You can think of PSRR and crosstalk in terms of impedance (usually measured to ground). If the impedance is low, any small parasitic coupling capacitors from the node to VDD or other noisy lines have little effect. (Imagine it as a voltage divider.) But if a node or output has high impedance, it's easy for interference

## RESONANCE REFRESHER

Unless you happen to work with switching regulators or RF front ends, chances are that you have forgotten all about inductors and resonant circuits. Don't fret: everything you need can be found in Figure S1.

For an inductor and capacitor in parallel, you add the admittance ( $1/\text{impedance}$ ) of the two elements together to get the admittance of the combination. But unlike resistors, the impedance and admittances of these elements are vector quantities, meaning that they have both magnitude and phase, and the phases of capacitors and inductors are exactly the opposite,  $180^\circ$  apart. This means that when added together, there will be some cancellation in the admittance. At some frequency, the magnitude of the admittance of the inductor will equal that of the capacitor. In that case, the cancellation is perfect, and the resulting combination has zero (or very low) admittance, which is the same as infinite (or very high) impedance. This is the parallel resonance shown in Figure S1. Such LC resonators are often referred to as *tank circuits* by analogy to some mechanical effect, not because they threaten to "tank" your project.

If there were no resistance in the circuit, there would be no theoretical limit on the impedance of a parallel LC combination. But there is always resistance (unless you are working at a few degrees Kelvin). The damping effect of the resistance is often expressed as the quality factor ( $Q$ ) of the resonance. There must be dozens of different formulas for  $Q$ , but it's easy to think of it as the ratio of peak impedance




FIGURE S1: Parallel resonance.

to the impedance that the inductor or the capacitor would have on its own at the resonant frequency. This latter quantity is sometimes called the *characteristic impedance* of the tank.

See. That wasn't so hard.

to disturb it. The self-correcting action of feedback makes loop nodes look like they have low impedance. However, as the loop gain drops, the impedance rises, making the nodes vulnerable to any source of coupling.

With Bandini Mountain choking off the supply, 100 MHz is going to be the noisiest frequency, but your feedback circuit is probably not going to be able to suppress it. For 60 dB of loop gain to be at full strength at 100 MHz, it would need to have a unity gain frequency of 100 GHz (assuming single-pole roll off). That's not going to happen. In fact, many practical circuits have already crossed unity gain by 100 MHz, so feedback provides no PSRR protection against Bandini Mountain at all. Your block cannot protect itself from supply junk, and it cannot hold its own output steady against coupling from nearby wires.

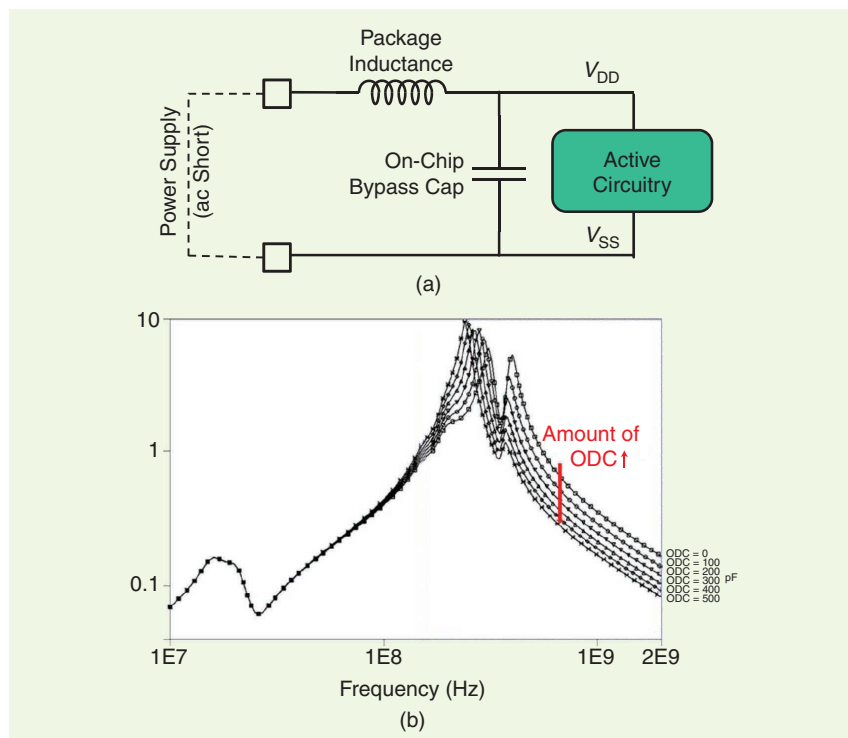
Figure 3 illustrates this effect with an example circuit. An idealized feedback loop with 60 dB of loop gain surrounds a vulnerable high-impedance node, labeled "X." This node is coupled to VDD by a 10-fF capacitor. To simulate the effect of PDN resonance, a "fake noise" generator uses a tank circuit to create a Bandini-shaped ac signal that is superimposed on the VDD supply.

Please note that the circuit in Figure 3 does not generate noise. It effectively uses a swept sine wave with an amplitude-versus-frequency profile similar to the dreaded Bandini Mountain so that we can see how noise might propagate through the example loop. Also note that you cannot reliably use this method to verify a real-life circuit because you cannot know the exact Bandini frequency and shape, nor can you accurately predict the total chip noise that will stimulate the Bandini tank. Figure 3 is just an illustration.

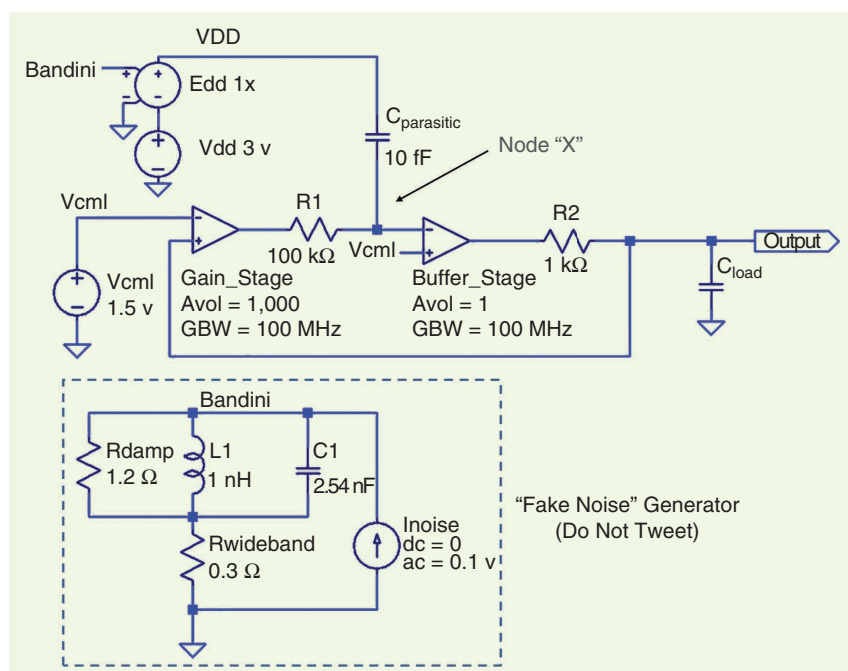
Figure 4 shows how the circuit of Figure 3 performs over frequency. ( $C_{load} = 0$  for Figure 4.) High loop gain keeps the impedance of "Node X" modest at dc. But as the loop gain declines in single-pole fashion, the "Node X" impedance rises at 20 dB/decade. At the same time, the impedance of the  $C_{parasitic}$  coupling capacitor is falling at

–20 dB/decade. So, the total coupling is increasing at 40 dB/decade, as is evident in the PSRR. By 100 MHz, the PSRR has dropped to 0 dB, and virtually all the Bandini noise makes its way to the circuit output.

I don't want to keep you awake at night, but it gets worse. Most feedback loops don't cross over 0 dB with a full 90° of phase margin. It's not unusual to have a modest amount of peaking in the closed loop response.



**FIGURE 2:** Your power supply from 50,000 ft. (a) The basic setup. The supply rails are usually considered to be an ac short at the board. (b) Bandini Mountain: impedance versus frequency in Hz as seen from the on-chip circuitry. ODC: on-die capacitance. (From [1]).



**FIGURE 3:** An example loop victim with Bandini-shaped noise on the VDD. GBW: gain-bandwidth.

Normally, that's not an issue, but that peaking means that the feedback is amplifying the coupled noise at that frequency. So, your feedback is providing gain—not suppression—to power supply noise at a frequency that is most likely in the

same frequency range as Bandini Mountain. Great.

### Threat #3: The Cap Gap

Out of concern for degraded high-frequency PSRR, or out of habit, designers often turn to passive filtering to protect

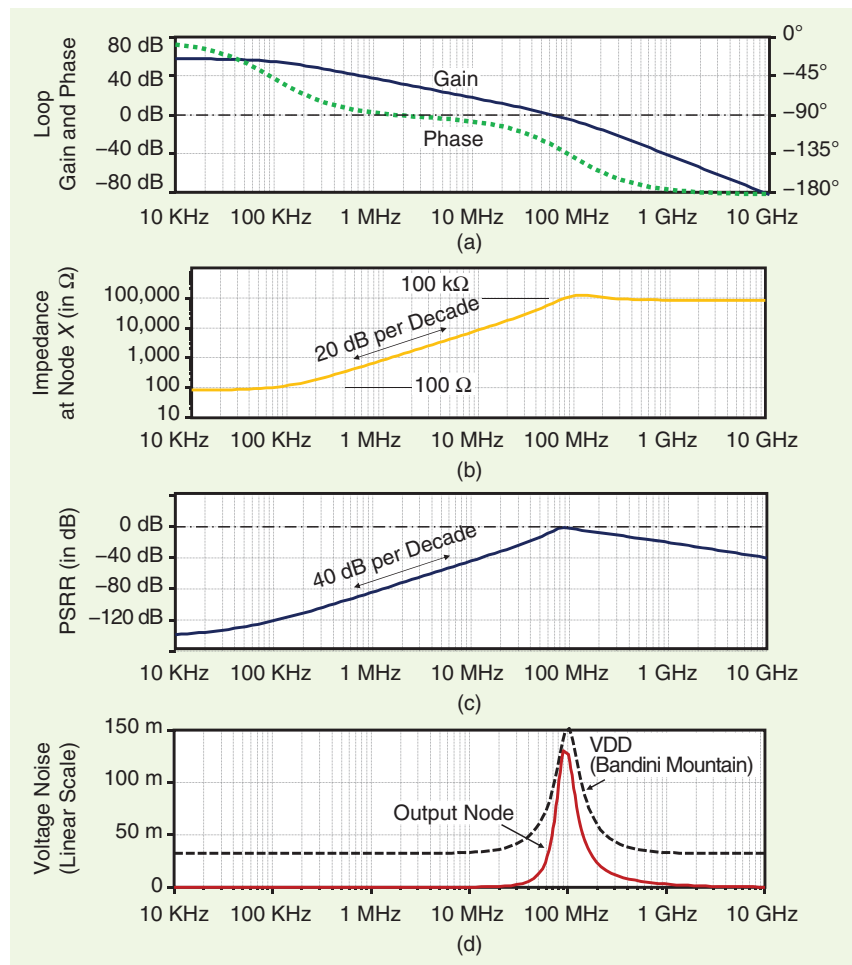
vulnerable low-frequency signals. They hang capacitors on sensitive lines with the assurance that high frequencies will be swallowed by their low impedance. And they will... if the capacitors are big enough. But the truth is that caps can rarely match the performance of feedback once you're in Bandini National Park. Sixty dB of feedback can make a 1-k $\Omega$  node impedance look like 1  $\Omega$ . To make a 1- $\Omega$  impedance at 100 MHz when the feedback gives up, you need nearly 2 nF of capacitance! Sure, in the multi-GHz range, a few picofarads are fine, but that doesn't cut it at 100 MHz... where you need it most.

So, if you plot the output impedance of a typical circuit versus frequency, you will see the impedance start to rise from 100 kHz to 1 MHz as the feedback loop gain drops. Eventually, the impedance starts to decline again as the filtering capacitors take over. But this leaves you with a peak in impedance, and that peak is—you guessed it—almost certainly in Bandini Land.


Figure 5 illustrates this for the example circuit of Figure 3, but there's not much to see. A modest filtering capacitance is added to the output node (see the blue line labeled 0.18 pF in Figure 5), but this light load doesn't even move the needle until 1 GHz, so it has zero effect on the Bandini noise. If you really wanted to change things at the Bandini frequencies, you'd need a capacitor several orders of magnitude bigger.

### Threat #4: The Patient Is the Disease

OK, let's say you don't care about chip real estate, and you're willing to



**FIGURE 4:** “Captain! The shields are down.” Loop gain provides no protection against Bandini Mountain for the example circuit in Figure 3. ( $C_{load} = 0$ .) (a) Open-loop gain and phase. (b) Impedance at node X without  $C_{parasitic}$ . (c) PSRR measured at output. (d) Voltage noise on VDD (dashed line) and voltage noise on output (solid line).



**FIGURE 5:** Futile attempts to filter Bandini noise with load capacitor. (See Figure 3.) (a) Output node impedance. (b) Noise on output node.



throw capacitance at the supply noise problem. So, you go ahead and start loading down the output node of Figure 3 with bigger and bigger capacitors. Success? No, disaster! You see, a node whose impedance increases with frequency is essentially behaving like an inductor. The last thing you want to do with an inductive node is to hang a capacitor on it. Figure 5 shows this with alarming clarity; see the red traces labeled 1.8 pF.

Let's get this straight. The loss of loop gain at a high frequency that makes nodes vulnerable to power supply noise also makes the same nodes inductive so they become the archenemies of the very capacitors that might save them from the supply noise. (For those new to the

planet, this is known as a *Catch-22*, "a frustrating situation in which one is trapped by contradictory regulations or conditions" [3].) Small capacitors can't protect vulnerable nodes, and large capacitors turn their wards into oscillators. And in case you're wondering, Goldilocks, there is no capacitor that is "just right." (Note: If you'd like to read more about this feedback-induced resonance phenomenon and how to analyze it, please refer to [4] in this series.)

As if a circuit resonance was not scary enough, a ringing circuit takes huge gulps of current, so it is possible to get a big spike in the supply current at the exact frequency range where you're most vulnerable. Since the Bandini phenomenon is a chip-

wide feature, it's unlikely that one unruly block is going to bring down the whole chip. But it can certainly do some damage locally if the offending circuit is supplied by a weak, narrow power plane.

## A Recipe for Trouble

Let's review.

- 1) Package inductance and on-chip capacitance *will* resonate to make your supply very noisy in the neighborhood of 100 MHz.
- 2) 100 MHz is beyond the bandwidth of many feedback loops, leaving them particularly vulnerable to this noise.
- 3) Small filtering caps have too much impedance to do any good at 100 MHz.

### THE 250-PS SPEED LIMIT


Parasitic resistance in capacitors makes for some interesting paradoxes. As you probably know, on-chip capacitors inherently carry a certain amount of resistance, known as the ESR ("equivalent series resistance") in discrete capacitors. While it certainly depends on wafer processing and some structural details, traditional 1-pF capacitors run about 250  $\Omega$ . That sounds like a lot, but it means your decoupling network should act like a capacitor until 640 MHz, well beyond Bandini territory.

At first glance, the ESR seems like a bad thing in decoupling caps because it causes extra voltage drop as the capacitor is trying to absorb fast transient currents. It turns out, however, that the ESR is *the* dominant damping mechanism for Bandini resonance. If the ESR were to magically go to zero, the Bandini peak would make its way to infinity. Be careful what you wish for. (Of course, there are other ways to damp a tank circuit. Please see [5].)

But there's even more counterintuitiveness to be had here. You might think that you could suppress the Bandini peak by adding more on-chip capacitance. You'd be wrong. Sure, the additional capacitance moves the resonant frequency (and that *might* be a good thing). But if you're trying to lower the peak impedance, you are in for a shock, and—again—it's because of the ESR.

You see, if a 1-pF cap effectively has 250  $\Omega$  in series with it, two such 1-pF caps in parallel would have a total of 125  $\Omega$  of the ESR. As the capacitance goes up, the ESR goes down. In fact, the time constant formed by the capacitance and the ESR remains the same, independent of value of the capacitance. That is,  $1 \text{ pF} \cdot 250 \Omega = 2 \text{ pF} \cdot 125 \Omega = 250 \text{ ps}$ . So, for any given capacitor structure in any given technology node, the self-discharge time constant is... well... a constant. And it limits the useful frequency range of the capacitor.

What this all means for Bandini is that although adding capacitance lowers the characteristic impedance (see "Resonance Refresher"), it also lowers the damping. Consequently—as luck would have it—the peak impedance doesn't change at all. It's not even a zero-sum game; the higher  $Q$  caused by the extra capacitance causes the tank to ring longer, which you probably don't want. This is illustrated in Figure S2, but you can see a real-life example in Figure 2(b), where adding additional capacitance actually makes the peak slightly worse. Weird, huh? Truth is stranger than fiction.



**FIGURE S2: Bandini Mountain for three different capacitances. The bigger the capacitor, the smaller the equivalent series resistance (ESR), and the higher the  $Q$ .**

4) Large filtering caps will destabilize your feedback loop.  
Want to rethink your career choice?

## Take a Deep Breath

It's standard editorial policy, even in mainstream news magazines, to capture the reader's attention with frightening headlines or alarming stories of gloom and doom but then to offer some potential solution or some small ray of hope so readers leave the story with a grateful sense of relief that maybe things will turn out all right after all. Not here. It's good for engineers to be anxious. Mother Nature and her evil brother Murphy *are* out to get you. You need to keep your guard up.

There is no easy fix for VDD problems. While we often get by with just dumb luck, these problems will not fix themselves. And chances are, they will eventually trip you up... if they haven't already.

That's not to say there is nothing you can do. For example

- Come up to speed on PDN issues. Reference [2] is an excellent source of practical info and is quite easy

to read. Large SoC teams will usually have somebody assigned to worry about this stuff. Make contact. There may be a supply model you can simulate with. Don't expect such models to be too precise, though. Resonant frequencies can only be estimated, even when the layout is known.

- Check the layout. Do the VDD or GND planes look "fingery" and inductive?
- People working at high speed are likely to take steps to damp their on-chip bypass caps. See [5] for an efficient solution. You may also want to take a look at "The 250-ps Speed Limit," because capacitors are not always what you expect.
- Good circuit practices can eliminate—or at least mitigate—many troublesome sources of supply coupling. If you'd like to see an article on this topic in this column, send your requests here: shifobrain@ieee.org. As always, be sure to include flattering comments about how much you love this series.
- Simulate PSRR, but remember that simulators think everything is

perfectly balanced. Real circuits are not. Imbalance your differential or pseudo-differential circuit with mismatch and signal swing.

- Understand how your feedback circuit can become inductive. See [4] for more on this topic.

OK, that's it for now. I leave you with the parting wisdom of Elvira, Mistress of the Dark whose tagline was "Unpleasant dreams!"

## References

- [1] G. Kim, D. Lim, T. Das, E. Lee, and S. You, "System Level Power Supply Induced Jitter Suppression for multi-lane high speed serial links," in *Proc. IEEE 72nd Electron. Compon. Technol. Conf. (ECTC)*, 2022, pp. 1863–1868, doi: 10.1109/ECTC51906.2022.00293.
- [2] L. D. Smith and E. Bogatin, *Principles of Power Integrity for PDN Design – Simplified*. Englewood Cliffs, NJ, USA: Prentice Hall, 2017.
- [3] "Catch-22." Dictionary. [Online]. Available: <https://www.dictionary.com/browse/catch-22>
- [4] C. Mangelsdorf, "Oscillation and ringing: Finding worst-case capacitive load [Shop Talk: What You Didn't Learn in School]," *IEEE Solid-State Circuits Mag.*, vol. 13, no. 2, pp. 16–20, Spring 2021, doi: 10.1109/MSSC.2021.3072801.
- [5] C. Mangelsdorf, "Duct Tape for the Frequency Domain: If This Ol' Hack Don't Fix It, It Ain't Broke [Shop Talk: What You Didn't Learn in School]," *IEEE Solid-State Circuits Mag.*, vol. 13, no. 4, pp. 15–22, Fall 2021, doi: 10.1109/MSSC.2021.3111427.


SSC

## THE ANALOG MIND (continued from p. 10)

### Complete Divider

We are now ready to create a chain of nine modules so as to meet the necessary divide ratio range. Note that the MC input of the ninth module must be tied to  $V_{DD}$ . The overall circuit draws a power less than 3 mW. Figure 13 plots the waveforms at nodes *E* and *F* of this module with  $f_{in} = 30$  GHz and for an overall divide ratio of 575. Let us examine the droop seen in these voltages.

The dynamic nature of  $C^2$ MOS logic makes it susceptible to leakage currents drawn by the transistors, an issue that becomes more serious for long clock periods. It also manifests itself to a greater extent at high temperatures, as the subthreshold leakage rises exponentially. In our design, the



**FIGURE 13:** The output waveforms of the complete divider exhibiting the effect of leakage.

last module must generate an output frequency equal to 50 MHz if the divider operates within a PLL sensing a 50-MHz reference. With the  $C^2$ MOS

latches residing in the store mode for 10 ns, their output states may degrade significantly. In the waveforms of Figure 13, the waveform at node *E* suffers from a droop of about 300 mV. Nonetheless, the inverters following this node restore the logical levels.

## References

- [1] B. Razavi, "The design of a millimeter-wave VCO [The Analog Mind]," *IEEE Solid-State Circuits Mag.*, vol. 14, no. 3, pp. 6–12, Summer 2022, doi: 10.1109/MSSC.2022.3184443.
- [2] C. S. Vaucher et al., "A family of low-power truly modular programmable dividers in standard 0.35- $\mu$ m CMOS technology," *IEEE J. Solid-State Circuits*, vol. 35, no. 7, pp. 1039–1045, Jul. 2000, doi: 10.1109/4.848214.
- [3] Y. Suzuki, K. Odagawa, and T. Abe, "Clocked CMOS calculator circuitry," *IEEE J. Solid-State Circuits*, vol. 8, no. 6, pp. 462–469, Dec. 1973, doi: 10.1109/JSSC.1973.1050440.

SSC