The Relentless Progression Of Power Supply Requirements

by Ken Coffman, Renesas Electronics America, Phoenix, Ariz.

It's not that I take everything personally, but it sure seems like the progression of power supply requirements is designed to drive me to distraction. Remember the good old days when 5 Vdc was common—with a wide tolerance of plus or minus 10%? All we had to do was deliver 4.5 V to 5.5 V with single- or double-digit currents. A fullscale volt of slop was acceptable!

Now, in every way that can be defined, power supply requirements are changing and getting harder. Why? Because modern loads are getting more and more complex. What are these modern loads? They are a nightmare of acronyms: CPUs, GPUs, FPGAs, ASICs, NPUs and AI/ML processors with stricter requirements and increased power needs every year.

Advertising and promotional material make it seem as if the task of creating a robust power supply design is trivial. To some extent, this is true. If the load is benign, then the challenge to the power supply is minimal, but that's the key point. Naturally, perfection is unattainable. But if the load requirements are unknown, then it is *impossible* to know if a design is good enough and good-enough gets harder every year.

What happens if your design *isn't* good-enough? If the control loop is out of compliance, the output can overshoot, undershoot or oscillate. If a power delivery system has resonances, then capacitors will be added to solve low frequency problems and create higher frequency problems. Any high impedance due to resonance—even at very high frequencies—can cause a supply voltage to be out of compliance—causing soft (recoverable), hard (recoverable only with a power cycle) or catastrophic (permanent damage) failures in the system.

What drives this madness? It's a combination of factors as I'll discuss in this article. I'll also explain why controlled impedance is the best tool for dealing with power supply requirements in real-world applications.

Smaller CMOS Process Geometries

As IC lithographies shrink, the allowable voltage across insulators also shrinks. A common main rail today is around 1.0 V and is headed lower. Suppose a modern chip requires a 0.72-V nominal voltage with a ±3% tolerance. That's ±22 mV or a range of 0.698 V to 0.742 V. This means the absolute accuracy must be good, but also, in addition to the nominal tolerance, the ripple and transient response must be very good too.

If given a very generous undershoot specification of -30%, then the output voltage can't droop, sag or undershoot by more than 216 mV, even for a short period. Tighter transient tolerances like -10%, -7% or even -5% are common. In some cases, an undershoot tolerance of zero is specified. This can be very difficult.

How much margin is needed for manufacturing and other tolerance build-ups? If the specification requires the load to never see anything less than -30% (-216 mV), the nominal design target should be something like 50%, or -100 mV. The purpose of a power supply is to make sure all the little transistors in the chip maintain their state and really act like transistors. They need a minimum V_{DS} to do that—which is represented by the core power requirements. Tables 1 and 2 show the nominal and initial tolerances you might see for a power supply rail of an FPGA.

Table 1. Typical core voltage requirements for three versions of Xilinx UltraScale FPGAs.[1]

	7 series (28 nm) V_{NOM}	UltraScale (20 nm) V_{NOM}	UltraScale (16 nm) V_{NOM}	UltraScale (16 nm) V_{LOW}
Operating voltage (VCCINT)	1 V	0.95 V	0.85 V	0.72 V

Table 2. Core supply voltage limits for -1 and -2 speed models of Xilinx UltraScale+ MPSoC FPGAs.

For -1LI and -2LE (VCCINT = 0.72 V) devices: PL internal supply voltage	0.698 V min.	0.720 V nom.	0.742 V max

© 2018 How2Power. All rights reserved. www.How2Power.com

As device geometries shrink, the power consumption per transistor decreases, but we are serving a lot more transistors, so the load currents are increasing. For perspective, an arc welder can emit 200 A to 400 A. There are modern loads requiring higher currents than this and the high-end requirements will soon go much higher. A tiny semiconductor device is not an arc welder—and does not require the higher voltage—but some semiconductors require *more* current than an arc welder. It's mind-boggling to consider.

Fig. 1 shows a volt-amp plot for a commercial welder.

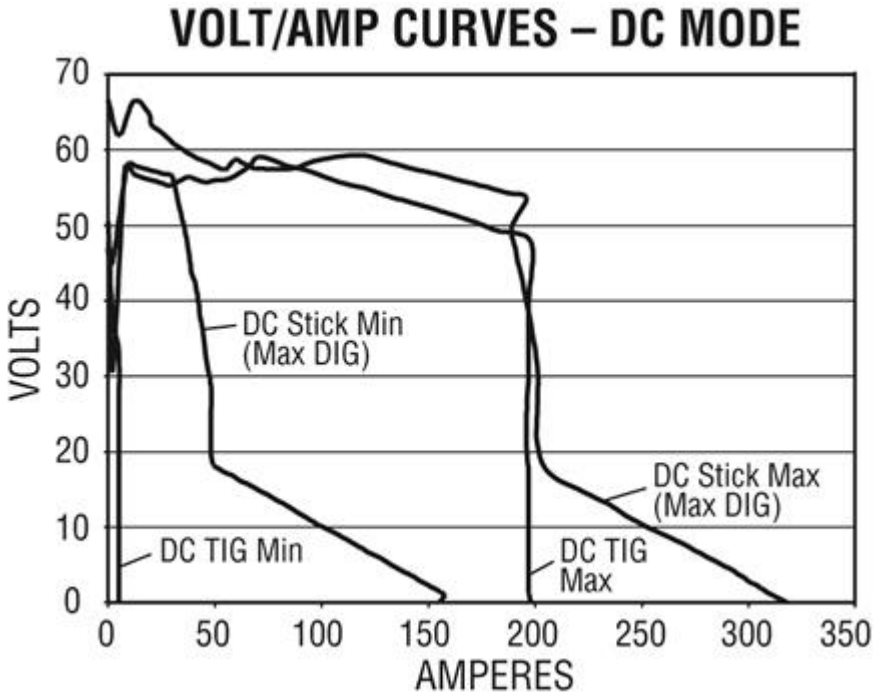

Fig. 1. Output current characteristic for an example arc welder. Graph courtesy of Miller Electric Manufacturing.[2]

Dynamically Changing Supply Voltage

Some power supply designs use AVS (Adaptive Voltage Scaling) to move the output voltage around depending on the operating mode. This means the output voltage for a stable regulator under load is changed. Is it just me, or is this like jumping out of a perfectly good aircraft? It would be one thing if we could take our time and get to the new operating point at our leisure, but that's not what is required.

The CPU wants to slew to a new operating point in a hurry—perhaps something like 25 mV/µsec. But from the power supply's point of view, we don't want to go too fast. First of all, we want to limit the peak current in the massive bank of output capacitors, but also most semiconductor devices have an ESD structure like an SCR on their power and I/O pins which, when tripped, turns on and will not turn off until power is removed.[3] Fast dv/dt transitions on power rails can inadvertently trip this ESD protection network—which is not what we want. So, what's a safe maximum dv/dt? 90 mV/µsec? 200 mV/µsec? 1 V/µsec? You tell me.

It would be great if the load currents were fixed, even if at a high current. Take an incandescent lamp, for example. After the filament heats up, the resistance is mostly constant. What a luxury!

However, to maximize the EVF (Engineer's Vexation Factor), modern electronic loads have slow changes (for example, between full-power and standby-power operating modes), fast changes (lots of registers changing states at some central switching frequency) and super-fast changes (for example: beat frequencies between different clock domains and harmonics from super-fast signal edge rates). It appears that anything that can cause problems for a power supply is mixed into the load profile. Thank you very much, modern technology.

© 2018 How2Power. All rights reserved. www.How2Power.com

The Importance Of Impedance

There are many ways to evaluate a power system, but I am convinced there is only one practical way to make sure a power supply design is stable, serves the load properly and does not cause problems for the clocks and other sensitive circuits elsewhere in a design. The power supply output impedance must be designed to be flat.

Istvan Novak, an authority on signal integrity and power integrity, has done some great work describing the "Big V" bypass capacitor bank, see Fig. 2. The green plot describes the Big V where a capacitor or capacitor bank, which could represent a bank of power supply bypass capacitors in a system, has a low-impedance resonance. Shown in Fig. 2, the Big-V impedance is flat up to 400 kHz, which is good. Then, the low ESR of the capacitor (or capacitor network) is shown with a self-resonant frequency (SRF) just north of 2 MHz; then the "excess" inductance kicks in and causes the impedance to increase.

With a low impedance resonance as shown, the high impedance at a higher frequency is inevitable. You can't "buy" the low impedance resonance at low frequencies without "paying" a horrible price at higher frequencies. A good power delivery network design will result in a flat impedance up to as high a frequency as possible. If the impedance is flat as shown in the blue plot, then the transient response, ringing and emissions will be as good as they can be. This idea is demonstrated by the step responses shown on the right in Fig. 2.

To explain Fig. 2, the lefthand plot is impedance, shown starting at 10 mΩ at 10 kHz. With those three impedance strategies, we see the corresponding step response on the right. From these plots, it is easy to see the dV damage created by uncontrolled resonance.

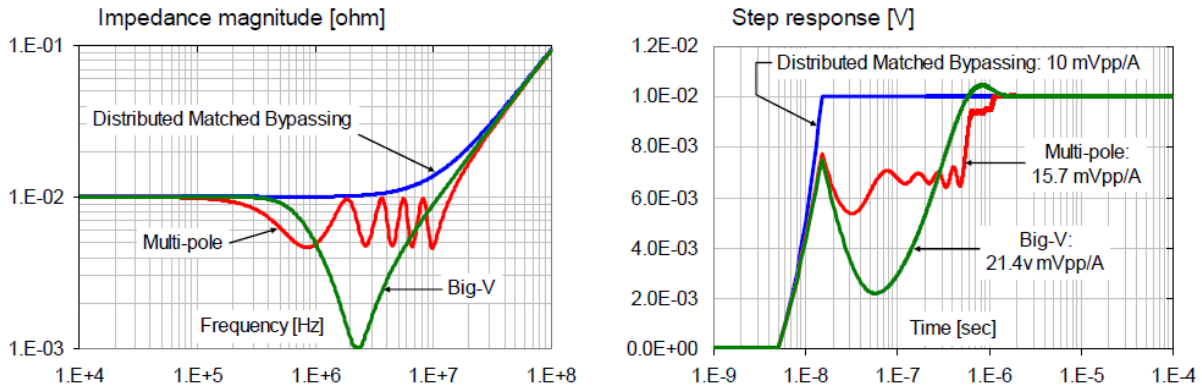

Fig. 2. Impedance of a Big-V capacitor bank (green curves) compared to multi-pole (red curves) and matched bypassing (blue). (Plot courtesy of Istvan Novak.[4])

The next question is: what impedance should be the target? That can be a very difficult question, but, absent a more thoughtful specification, Larry Smith and Eric Bogatin suggest using the maximum power dissipation of the IC package. Table 3 offers an example.

Table 3. Using IC power dissipation limit to calculate target impedance.[5]

Segment	TDP (W)	C1E (W)	C6 (W)
1.3 mΩ R_{LL}	85 W	--	15 W

What load current does this thermal design power (TDP) limit represent?

$$i = \frac{P}{v} = \frac{85}{0.72} = 118Apk$$

What is a reasonable delta-current (di)? How about a 50% or a 60-A step?

Using our example of 100 mV of transient variation, we can calculate a target impedance.

$$Z = \frac{dv}{di} = \frac{0.1}{60} = 1.67mOhms$$

© 2018 How2Power. All rights reserved. www.How2Power.com

This impedance target is undoubtedly very conservative. What's wrong with that? If the design team complains, they can provide a more practical power dissipation based on what the part is expected to dissipate in the customer design. For this example, we want the PDN impedance to always be less than 1.67 mΩ, but not a lot less. If the target impedance is greater, that makes the design task easier. Don't make this decision on your own, only relax the requirement based on what the IC design team tells you.

What does this impedance mean? If all the transistors in the ASIC draw 85 W and our source impedance is greater than 1.67 mΩ, then the supply voltage will drop by more than 100 mV. It doesn't matter what frequency this happens at. The modern load has a very wide bandwidth and the worst-case frequency will be stimulated at some operating condition, guaranteed.

For the elements under our control: the control loop if we're designing a regulator system, the PWB if we're laying out a board or the ASIC interconnect if we're designing the chip, we must have a controlled impedance not higher and not much lower than the target.

Why not much lower? As already stated, for each low-impedance resonance, a high impedance price will be paid at a higher frequency. Did you expect to get something for nothing?

At a glance, it seems like the target impedance design strategy imposes unwanted complexities for us. But looked at another way, we should be thankful. We are all searching for relevance and longevity in a fast-changing world. Rock-solid power supply designs are not going to just spontaneously emerge from ignorance. Engineering rigor and designers who think about these issues are required. That's good!

Preparing For The Challenges

As responsible power supply designers, what are we going to do to adapt to upcoming challenges? First, there is a required reading list, here it is:

- *Power Integrity: Measuring, Optimizing, and Troubleshooting Power Related Parameters in Electronics Systems* by Steven M. Sandler
- *Principles of Power Integrity for PDN Design—Simplified: Robust and Cost Effective Design for High Speed Digital Products* by Larry Smith and Eric Bogatin.

Next, partner with your preferred semiconductor vendor to create systems that are well-controlled and efficient. How many bypass capacitors do you need? That depends on the controller bandwidth. Can you use only MLCC or do you need POS or SP caps too? If the controller is stable and responds quickly with a high crossover frequency, then the requirements of the capacitor bank are reduced.

Take a look at the available power supply controller IC offerings such as Renesas Digital Multi-Phase product offerings.[6] But whichever power supply IC manufacturer you choose to work with, challenge them and see what they can do to help.

I work with many engineers who understand the requirements and tradeoffs for creating a well-behaved power supply better than I do. At the same time, there are many who don't know what they don't know—who are thrust into design situations before they are ready. If you study your load and understand its requirements, you are on the path to success. If you can't get solid information about what your load requires, then assume the worst, document your assumptions and be ready to change your design if the assumptions fall short of what is needed.

Good luck out there, my friends.

Acknowledgements

The author thanks Steve Sandler, Tammie Bard, Istvan Novak, Eric Bogatin, Chris Thompson, Hunter Chou, Preeti Khanna and Keven Wei who reviewed the manuscript drafts, highlighted errors and made helpful suggestions.

References

1. Operating voltages and 0.72-V specification courtesy of Xilinx, https://www.xilinx.com/support/documentation/data_sheets/ds925-zynq-ultrascale-plus.pdf.
2. Arc welder volt-amp curves courtesy of Miller Electric Manufacturing, https://www.millerwelds.com/resources/article-library/selecting-a-constant-current-cc-dc-welder-for-training-purposes.

© 2018 How2Power. All rights reserved. www.How2Power.com

3. For comments about ESD structures, see: https://e2e.ti.com/blogs_/b/analogwire/archive/2015/12/02/what-you-need-to-know-about-internal-esd-protection-on-integrated-circuits.
4. One of Istvan Novak's discussions about Big V capacitor networks is here (this is the source for Fig. 2): http://electrical-integrity.com/Paper_download_files/DC06_TF-MP3_SUN.pdf.
5. TDP specification courtesy of Intel.
6. Intersil's Digital Multiphase Controllers page.

About The Author

Ken Coffman, based in Phoenix, Arizona, is a Senior, Principal Field Applications Engineer with Renesas Electronics America. Ken has been working on multi-phase power supplies for 15 years and therefore, considers himself a novice who is still learning. He can be reached at Kenneth.Coffman.xh@Renesas.com

Do you have any comments on this column? If so, share them with the editor.

© 2018 How2Power. All rights reserved. www.How2Power.com