Making a Steamy, Hairy Golf Ball

Istvan Novak, Samtec, Life Fellow of IEEE

In his interesting article [1], my friend Steve pointed out a major hurdle we face in power distribution design: power engineers (who design power converters) and power integrity engineers (who design system bypassing-decoupling networks) use different vocabulary, techniques and requirements. To understand a little better how we got here, I want to start with a prediction I heard sometime in the early 90s at one of the conference keynote speeches: "in ten to twenty years, computers will look like hairy steamy golf balls" The argument was that the many different IO cables (do you remember parallel printer cables?) will make computers look hairy and because they will be much smaller, probably the size of a golf ball, and they will dissipate more power, they will be steaming hot. Almost thirty years later it is time to ask: are we there yet?

Going back a few more decades, *Figure 1* shows a computer board made of germanium diodes and transistors.

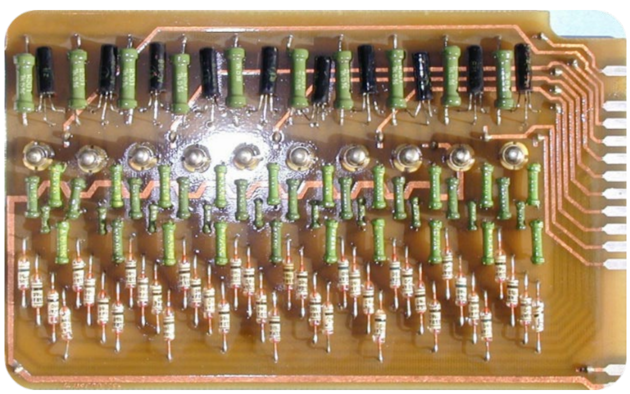

Figure 1: *A DTL (Diode-Transistor-Logic) from the early days of computers. You can notice that there is no power distribution component whatsoever on the board.*

This computer board was working fine with no bypassing or decoupling component on the board. There was a central power supply and a few bulk capacitors in the chassis, that was all about power distribution. In contrast, today the power-hungry big chips need hundreds of power distribution components to feed them, and those components fight for the same prime space around the chips that critical signal-integrity and thermal solutions also want to occupy. *Figure 2* shows a small personal computer motherboard from recent years.

Figure 2: *Partial top view of a Gigabyte H110M-A motherboard.*

Practically all of the components you see around the CPU are serving various power distribution purposes.

What makes power distribution design really challenging today is illustrated by the pyramid of requirements shown in *Figure 3*. On the top, we have the system functionality requirements. The customer ultimately cares for the system functionality, availability and reliability (along with cost, of course), and therefore this is the level where the system designer and manufacturer have a lot of detailed specifications and requirements to adhere to. These requirements may be challenging and may be hard to achieve, nevertheless they provide guidance for the design and validation. Also, determining whether the system meets the requirements is (relatively) simple and straightforward. This all changes, however, as we go down in the pyramid to look at the supporting functions. For high-speed computing and communications boxes, signal integrity is unquestionably a fundamentally important support function. Signal integrity guarantees that the bits get transferred with no or just with a sufficiently few errors.

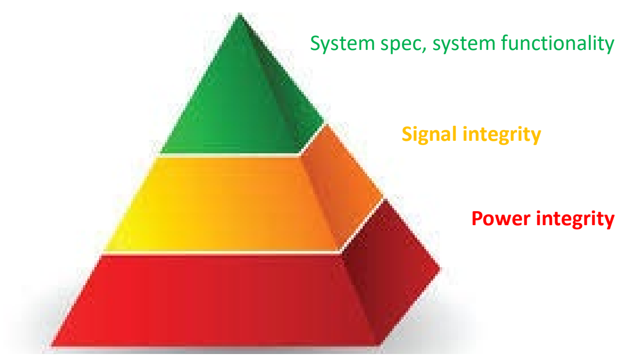

Figure 3: *Pyramid of various requirements for a system.*

Gradually over the decades, industry-wide standards, best practices and procedures have been developed to help the signal-integrity design, bring up and validation. We have electrical requirements spelled out in the physical layer section of all signaling standards. These requirements have come a long way: first we had only bit error rate requirements, later various masks came along, and today we have Effective Return loss (ERL) and Channel Operating

Margin (COM) helping the design and validation. Test instruments and simulation software alike incorporate these requirements to help our test and validation. Also, several signaling standards offer fail-safe fallback: if the high-speed link does not come up or stop working at the maximum speed, it scales back the data rate until at least some functionality is restored. So my point is that in signal integrity, today there are a lot of specifications, requirements and standards providing guidance and setting design goals.

The availability of support and guidance changes dramatically as we move further down on the pyramid. I show power integrity at the bottom of this pyramid, for two reasons. First, power integrity is really the necessary foundation for the layers above, for signal integrity and system functionality. If power distribution does not work properly, signal integrity and system functionality may become totally irrelevant; after all the system will not work. Second, for power integrity designs and validation there are no industry standards and often even the specifications and requirements are awfully incomplete. Along the power distribution design process, a lot of the requirements are left to the designer to establish, what makes the design and validation processes a little bit ill-defined and very challenging.

As we circle back to the prediction that was made in the 90s: you can imagine how the power distribution looks like for power-hungry chips requiring DC-DC converters with up to 1000A of maximum current [2]? Though a few years late compared to the prediction, large switch chips handling many terabytes of traffic soon very well may look like hairy, steamy golf balls.

And now if we finally return to the challenge pointed out in [1]: why power engineers and power-integrity engineers still have differing vocabulary and different approaches to the same common problem? *Figure 4* will help us to explain.

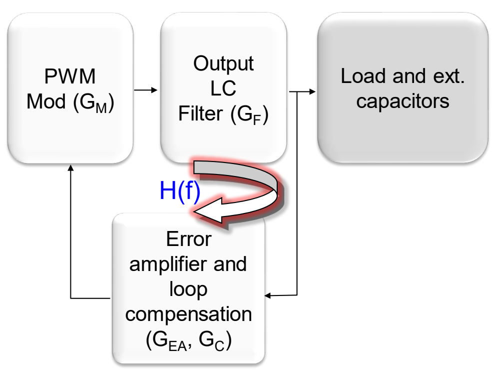

Figure 4: *Block diagram of a switching voltage regulator.*

The white blocks in the figure represent the main subsystems of a switching regulator: this is the territory of the power engineer. The gray block (the territory of the power integrity engineer) represents the circuit in the user's system: bypass capacitors, optional filters, and eventually the load(s) we need to feed. These blocks are connected together at the converter output. The output voltage is regulated by a feedback loop and the sense point of the feedback circuit has to monitor the same output where the converter circuit and the user load connect. This means that from a feedback-loop stability point of view the voltage regulator circuit and the user load are inseparably linked together and ultimately the stability can be evaluated only if we know all of these blocks. When we design a DC-DC converter for a given, known load, we can do the power and power-integrity designs in close collaboration or doing both by the same people. But when power engineers design voltage regulators for the open market, how could they know the user requirements and user circuits? Of course, they don't. Also, for reasons of economy, voltage regulators tend to have wide input and output voltage ranges, further increasing the vast number of possible user applications. All of the above left the power engineers with little choice, but to focus on the regulator itself and isolate it as much from the load as possible. During the early decades of switching regulators this resulted in limited loop bandwidth and a substantial bank of output capacitors, such that the impact of adding the unknown user circuit was minimized.

The DC-DC converter shown in *Figure 5* is a twenty-year old design and supplied a maximum of 16A. The useful area of the converter board (without the edge fingers) is 1.25" x 3.5". Next to the output connections we see sixteen large ceramic capacitors. These capacitors, together with the series inductance of the card-edge connection provided substantial isolation from the unknown load impedance. In those days the fact that power engineers and power-integrity engineers spoke different languages, did not create a significant problem yet. Today, as a result of improved technology and ongoing cost reduction, a dual 9A (or single 18A) fully encapsulated DC-DC converter module has a single 2.2uF capacitor on each of its outputs and comes in a 0.63" x 0.5" BGA package [3]. The BGA package provides low connection impedance, but also less series isolation. Similarly, the single ceramic capacitor on the output means that to analyze the stability of the feedback loop, we can no longer ignore the interaction between the power converter output stage and the load circuit. It is now time to unify the vocabulary and design/validation processes of the power and power-integrity engineers.

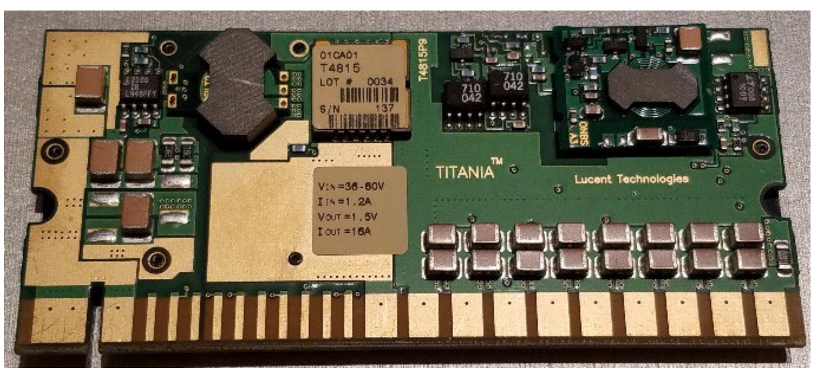

Figure 5: *A Lucent Technologies 16A DC-DC converter.*

At the end you may ask how all this relates to the hairy and steamy golf ball? Well, the same underlying trend of miniaturization, which pushes the power density to the boiling point, removes the isolations and barriers among previously separate system domains, forcing us to rethink how we design power converters and system power distribution.

References:

[1] Steve Sandler, "Power Electronics vs. Power Integrity," Signal Integrity Journal Print Edition, January 18, 2019

[2] https://www.infineon.com/cms/en/product/promopages/Powering_Next_Gen_Processors/

[3] https://www.analog.com/media/en/technical-documentation/data-sheets/4675fb.pdf