Power Distribution Network (PDN) Impedance and Target Impedance

Steve Sandler, Picotest, Eric Bogatin, Teledyne LeCroy, Larry Smith, PDNpowerintegrity.com

Introduction

The purpose of the power distribution network, the PDN, is to deliver a DC voltage, within regulation limits, with an acceptable noise, to each active device. Ultimately, the final design is a balance between managing the acceptable performance with sufficient margin and an acceptable level of risk given the ever-shrinking schedule and cost target constraints.

When starting out with a design, using a target impedance, as the maximum impedance design goal for the PDN impedance, is a useful metric to guide the design. But, just using a target impedance alone as the design requirement is not enough. There are a few significant noise sources from the PDN elements that a simple target impedance does not adequately address. To be clear, the voltage regulation limits that must be maintained include ripple, noise, DC regulation, IR drop, and the voltage excursions due to dynamic transients. For this paper, we are only discussing the voltage excursions due to dynamic current transients as the 'noise source' that is being managed by the PDN impedance control.

The Target Impedance

A simplified model of the PDN including the power generator element, or voltage regulator module (VRM), the power consumer elements, such as the IC devices, and the passive interconnects, including the various decoupling capacitors, is shown in Figure 1. This model can be used for both the I/O power rails, often labeled as the Vcc rails and the core logic rails, often labeled as the Vdd rails.

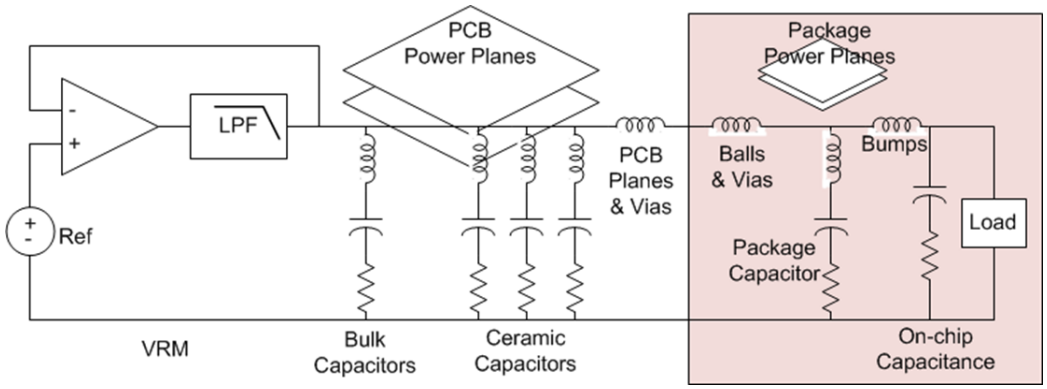

Figure 1 A simplified model of the PDN as seen by the pads of the die power rail.

There can be multiple sources of noise on the power rails of the IC, but separated into two general categories:

- **self-aggression noise** caused by the chip's own transient currents passing through the impedance of the PDN
- **mutual aggression noise**, or crosstalk, from the PDN interconnects onto the voltage rails of the IC. This pollution noise can come from the VRM, from other power rails, or even

from signals passing through the impedance of the PDN interconnects, as part of the return path of the signals.

Target impedance is a useful metric to address the first problem, in a few situations, but does not address directly the second problem.

When analyzing the self-aggression noise from the Vdd rail transient current on the Vdd rail, we consider the voltage noise generated on the PDN interconnects due to the transient currents from the core logic switching. The transient core logic power current through the PDN impedance creates the voltage noise on the core logic power rails.

To keep the voltage-drop from the transient current below the acceptable voltage noise threshold, the impedance of the PDN must be below a certain level, the target impedance. In its simplest form the target impedance is given by:

$$Z_{target} = \frac{\Delta V_{max-noise}}{I_{transient}}$$

If the impedance profile of the PDN, as seen by the IC pads, is a flat impedance profile from DC to the highest frequency components of the transient current, and we keep the flat PDN impedance below the target impedance, then the worst-case voltage noise, due to dynamic current, on the pads of the IC rail will be less than the maximum acceptable voltage noise.

An important PDN design goal is to engineer a flat impedance profile for the PDN, below the target impedance. This is the basic criterion for a 'robust' PDN. In this case, the worst-case voltage noise, due to dynamic current excitation, for a maximum transient current will always be less than the maximum acceptable noise.

While a flat impedance profile is an important goal, in many cases, a flat impedance profile across the entire frequency range is too expensive to implement. There may be peaks in the impedance profile. These usually arise at the interfaces to the boundaries between the different interconnect elements, such as

- between the VRM and bulk decoupling capacitors
- between the bulk decoupling capacitors and the MLCC capacitors
- between the package lead inductance and the on-die capacitance

Figure 2 shows an example of a typical impedance profile of the PDN, as seen by the pads of the Vdd rail showing a few peaks, some larger than others.

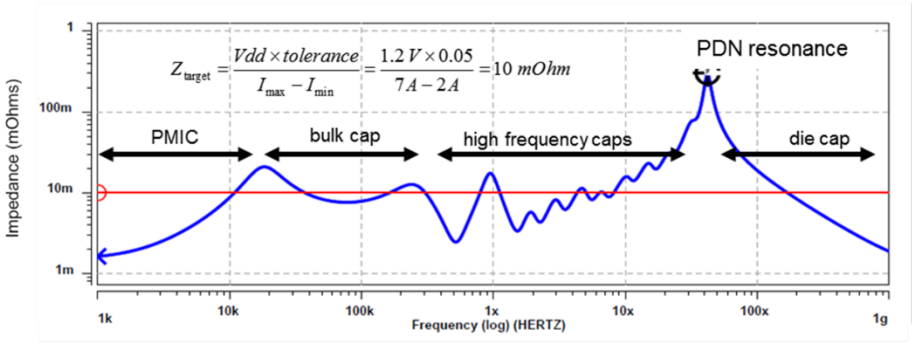

Figure 2. An example of the impedance profile of the PDN as viewed from the pads on the die power rail.

Using a target impedance as a design goal, is a robust solution in the specific case of a flat impedance profile. When there are peaks in the impedance profile, the worst-case voltage noise is not accurately predicted by just a target impedance.

To engineer the peak impedances below the target impedance may make the PDN very expensive. This would require many MLCC capacitors, and thin dielectric between the power and ground planes, for example. Instead, a second criterion is introduced for acceptable performance where peaks occur.

The Features of an Impedance Peak in the PDN

Every impedance peak in the PDN is created by a parallel RLC circuit. Four terms characterize each peak:

- The parallel resonant frequency
- The characteristic impedance
- The peak impedance
- The q-factor

The parallel resonant frequency is the frequency at which the inductive reactance equals the capacitive reactance. This is calculated from

$$f_{res} = \frac{1}{2\pi\sqrt{LC}}$$

The characteristic impedance is an impedance that characterizes the RLC circuit. It is not to be confused with the characteristic impedance of a transmission line but is just another impedance that characterizes the RLC circuit. This characteristic impedance is the impedance at which the reactive impedance of the capacitor equals the reactive impedance of the inductor, which occurs at the parallel resonant frequency. The impedance where they cross, the characteristic impedance, is given by

$$Z_0 = \sqrt{\frac{L}{C}}$$

The q-factor of the peak is related to the damping created by any equivalent series resistance, the R term, from lead resistance, dielectric loss or other interconnect resistance, and is given by

$$q-\text{factor} = \frac{1}{R} Z_0 = \frac{1}{R}\sqrt{\frac{L}{C}}$$

In the case of a q-factor greater than 1, the peak impedance is related to each of these terms by

$$Z_{peak} = q-\text{factor} \times Z_0 = \frac{1}{R}\frac{L}{C}$$

The importance of these terms is seen when analyzing the voltage across the capacitor from two important current waveforms, a step response and a square wave response.

Step Response of an Impedance Peak

As a first order model, the current draw from a die can be modeled as a current source. The bandwidth of a step current change, for a specific rise time is given by

$$BW = \frac{0.35}{\text{RiseTime}}$$

When the current source has a step current change with a rise time shorter than 3.5 nsec, the bandwidth is greater than 100 MHz. If the parallel resonant frequency is below 100 MHz, then the voltage response of the RLC circuit to a step current change is a damped sine wave. The frequency of the damped sine wave is the parallel resonant frequency. The peak amplitude of the damped sine wave is given by

$$V_{peak} = I_{max} \times Z_0$$

An example of this response is shown in Figure 3. This suggests when there is a step current change, to keep the peak voltage noise on the PDN below the acceptable noise voltage, the characteristic impedance of each peak should be below the target impedance. If this is the case, then the worst-case voltage noise on the PDN will be less than the maximum acceptable level. This is the next most restrictive criterion for a robust PDN.

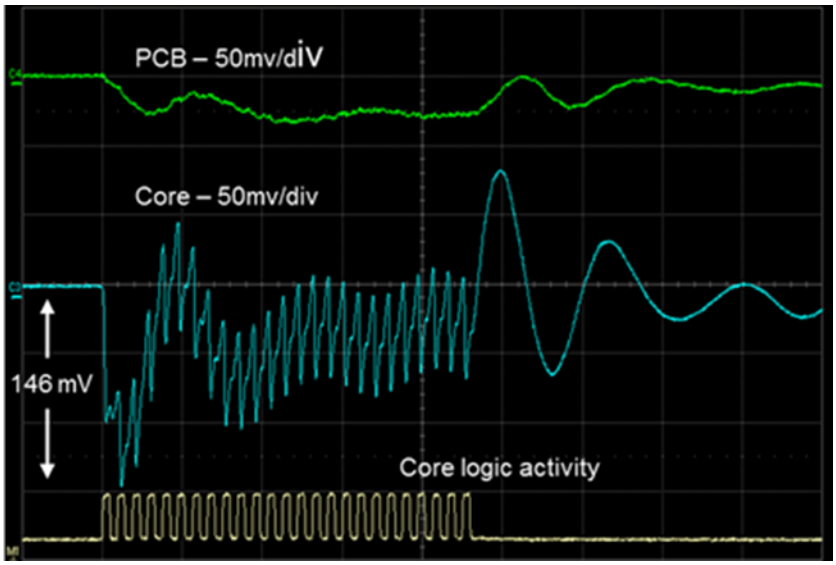

Figure 3. An example of the measured on-die voltage noise of a power rail as the core logic current undergoes a step change in current load. Superimposed on the step current are the individual clock edge noise responses. Note the damped sine wave during the release.

Often, engineering the characteristic impedance to be below the target impedance is also an expensive choice. Some risk might be acceptable that either there will never be a worse case transient current with a short enough rise time to excite the full peak, or this peak voltage noise might be acceptable if it is only present for a short period of time.

But another current waveform may generate even more noise from a peak impedance in the PDN.

Resonant Response

When the transient current is in the shape of a square wave, with a repeat frequency exactly at the resonant peak frequency, the peak to peak voltage noise generated by the sine wave frequency component of the square wave of current at the parallel resonant frequency is

$$V_{pk-pk} = I_{transient} \times \frac{4}{\pi} \times Z_{peak} = I_{transient} \times \frac{4}{\pi} \times q - factor \times Z_0$$

If the voltage on the die rail is at the nominal value, then this sine wave resonant response will be centered at the nominal value minus some IR drop. The noise excursion will be just the amplitude, which is half of this peak to peak value. The voltage above and below the nominal is:

$$V_{excursion} = \frac{1}{2} I_{transient} \times \frac{4}{\pi} \times q - factor \times Z_0 = 0.64 \times I_{transient} \times q - factor \times Z_0$$

For the same transient peak to peak current, but at the resonant frequency, the resulting voltage noise excursion can be 0.64 x the q-factor times the step response voltage noise, both above and below nominal voltage. In some critical mission industries, the maximum voltage is important because the circuits are not allowed to be electrically overstressed. In other industries, the minimum

voltage is important because set-up time requirements, Vmin and Fmax considerations. It takes about a q-factor of cycles for the resonant square wave of current to drive the voltage noise to the peak voltage. A q-factor of 1/0.64 = 1.57 or more will result in a peak noise voltage exceeding the maximum allowable noise when the Z_0 characteristic impedance meets the target impedance. But, to meet this criterion, not only must the maximum worst case transient current modulate at the parallel resonant frequency, but it must be present for a q-factor of cycles to reach this worst case.

The larger the q-factor of the parallel resonance, i.e., the sharper the peak due to smaller damping, the larger the resonant response can be, but the longer the PDN must be driven to reach the full peak noise. Figure 4 is an example of the measured resonant response voltage on the die when a square wave of transient current drives the on-die PDN.

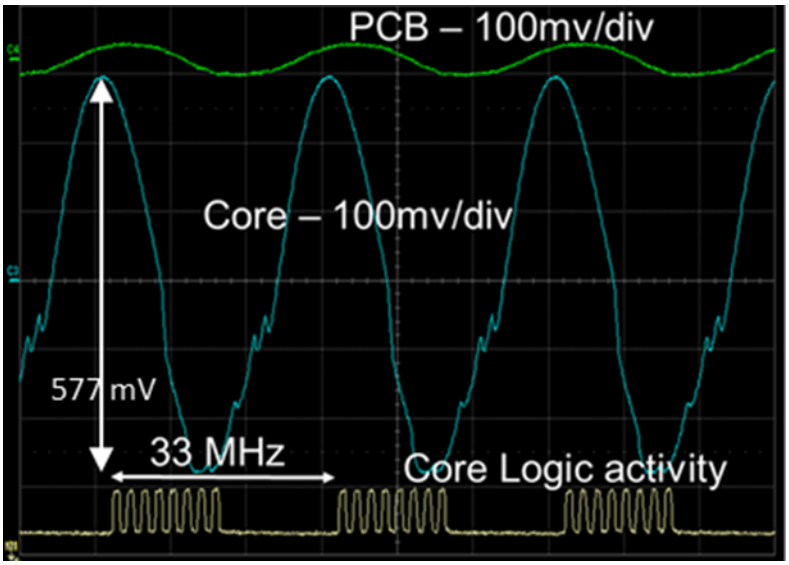

Figure 4. Measured voltage noise on a die when the core logic current is a square wave at the peak impedance frequency.

In a PDN designed to be fully robust and immune to resonant peak excitation, so the worst-case voltage is below the acceptable noise threshold, the parallel resonant peak impedance must be below 1.57 x target impedance. The 1.57 factor comes because the positive and negative resonant excursions are $0.64 \times I_{transient} \times Z_{peak}$ and the full voltage tolerance is reached at resonant conditions when an impedance peak is 1/0.64 = 1.57 x Z_{target}. The whole resonant waveform is biased and reduced by DC IR drop so the voltage minimums may be farther away from nominal voltage than the voltage maximums. Reducing the impedance peak below this level will keep the worst-case peak-to-peak noise below the acceptable level when the resonance is fully stimulated. But this is often very expensive, especially when the resonant peak is from the on-die capacitance and package lead inductance.

Often, some risk is accepted based on the unlikely possibility of the worst-case transient current occurring exactly at the parallel resonant frequency and lasting for a q-factor of cycles. Many successfully deployed products have a peak height as large as 3x the target impedance. This is an impedance peak height 2x higher than a fully robust PDN.

But even a fully robust PDN, with all of its peaks below 1.56x the target impedance, may still have PDN voltage noise above the acceptable limit under a special set of conditions. This is when rogue waves may develop.

Rogue Waves

If the PDN impedance profile has two or more high q-factor parallel resonant peaks, each of which is below the target impedance, none of the peaks, if excited individually, might generate excessive voltage noise under worst case conditions.

However, it is possible for a specific combination of resonant responses to sequentially build up exactly in phase to excite the lowest frequency peak, and then at its peak response, the second highest frequency peak, and at its peak, the highest frequency peak.

When multiple-frequency ocean waves add up exactly in phase, they can produce exceptionally large wave heights, which have been termed rogue waves. When multiple current waveforms are sequenced to excite multiple high q-factor peaks, the voltage noise on PDNs can show large peaks, which have been termed rogue waves.

Figure 5 is an example of the impedance profile of a PDN with three high q-factor parallel resonances.

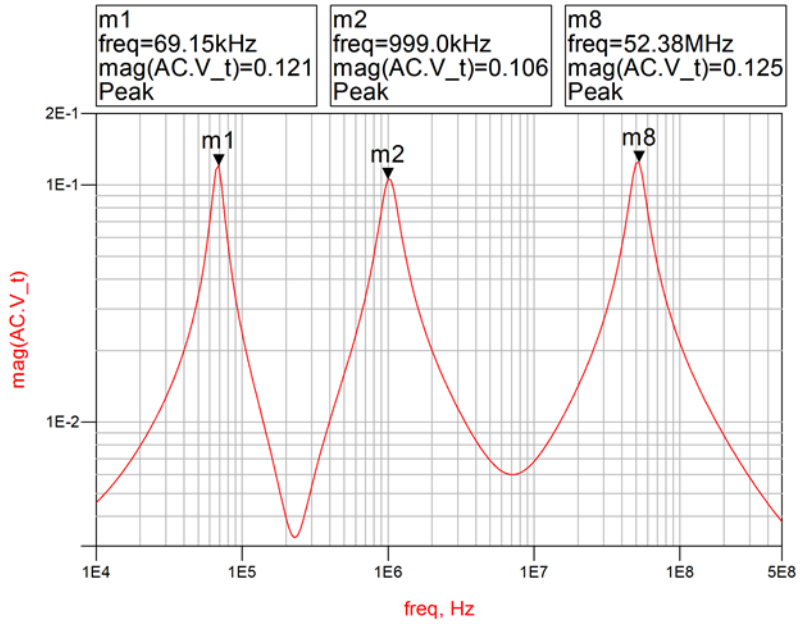

Figure 5. Example of an impedance profile with three, high q-factor parallel resonant peaks. These occur at the boundaries of various PDN elements.

When this PDN is excited by at least two sequential groups of square waves of current, with a duration and phase exactly optimized, they can excite rogue waves on the PDN. Each square wave excites a different resonance and the peak voltages build on each other. Even if each peak were below 1.56x the target impedance, meeting the fully robust requirement for acceptable noise, the resulting

rogue wave of voltage can easily exceed the acceptable level of noise. Figure 6 is an example of the simulated rogue wave on this PDN.

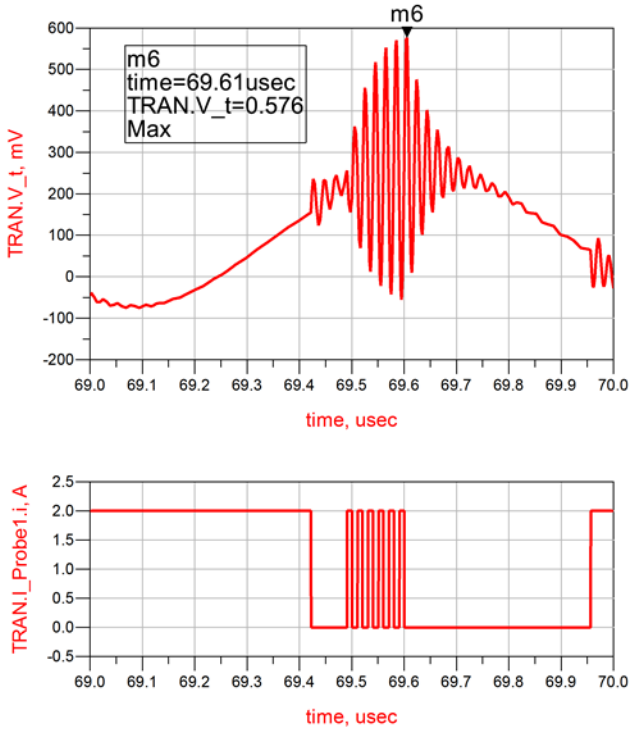

Figure 6. Example of the rogue wave created by sequentially exciting two of the resonant peaks, by the current waveform shown in the lower graph.

How likely is it that the transient current waveforms from the core logic circuits would exactly match the specific resonant frequencies and stay on resonance long enough to excite the peak amplitude perfectly phased? In a critical application, it only must happen once, causing the voltage requirements to be exceeded. This can result in an overstress condition, or intermittent operation. Figure 7 is an example of the measured rogue wave response on a PDN from the unfortunate occurrence of exactly the right current waveform.

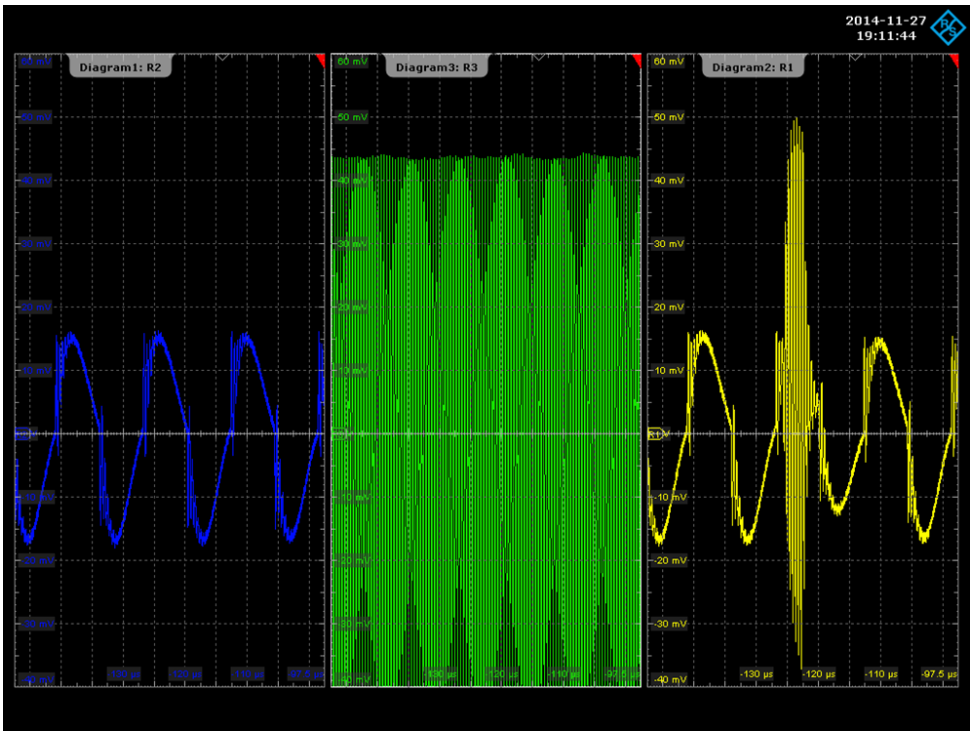

Figure 7. An example of the measured voltage noise on a PDN from two different square waves of current with the same amplitude, exciting each resonance. In the far-right screen, these two excitations are phased exactly right to excite the rogue wave with a response larger than either of the two other responses.

Other Sources of Noise

This analysis has so far only considered one source of noise, the self-aggression noise from transient currents on the Vdd rail through the PDN impedance profile. Target impedance is a useful metric against which to judge the impedance profile, if it is applied with the appropriate balance of performance, risk and cost. But there are also other sources of noise which can be uncorrelated and add to the Vdd noise.

Noise on the Vdd rail, for example, can come from self-induced noise and mutual sources from the noise pollution on the board level interconnects. The board level noise can arise from the VRM, from the transient currents from other consumer ICs and even from signals transitioning through the power rails. This noise is independent of a target impedance guideline. In the worst case, all these noise sources will add together with the self-aggression noise on the Vdd rails. A fully robust PDN should consider this worst case noise.

Conclusion

A robust PDN is engineered to have a flat impedance profile with an impedance below the target impedance. This should always be the design goal in every PDN. A flat impedance profile will eliminate the chance of resonant responses, rogue waves and provide a deterministic noise from the self-aggression Vdd currents.

When the impedance profile is not flat, but has parallel resonant peaks, keeping the peak impedances below the target impedance is not enough of a criterion for a fully robust PDN.

Many designs are a balance between the performance, the risk and the cost. When cost is an important driving force, an acceptable risk might be when a parallel resonance's characteristic impedance is engineered to be below the target impedance. This would keep the peak voltage noise response from a step current to be within the acceptable noise limit.

But even this criterion is not robust enough for rare current waveforms. When resonant driving of peaks is possible, all impedance peak heights should be below 1.56 x the target impedance.

With multiple high q-factor peaks, even this condition is not enough for a fully robust PDN. Under rare conditions, a rogue wave with a voltage amplitude above the acceptable voltage noise limits may be created when multiple peaks are sequentially excited in phase.

Finally, a fully robust PDN would need to include the worst-case combination of all noise sources on each rail, from self-aggression noise and mutual-aggression noise.

All designs are a tradeoff between the performance, the cost the risk and meeting the schedule targets. Each type of product has a different appropriate balance between these factors. In many consumer products cost is the biggest driving factor and some risk is accepted. In high reliability products, low risk is the driving force and all of these effects must be included in any PDN analysis. The target impedance is a starting place to guide the design in each design, but in low risk products not enough of a criterion to base a design.

References

1) Larry D. Smith, Eric Bogatin, "Principles of Power Integrity for PDN Design--Simplified: Robust and Cost Effective Design for High Speed Digital Products" (Prentice Hall Modern Semiconductor Design), ISBN-13: 978-0132735551, April 6, 2017
2) Sandler, S, "Power Integrity: Measuring, Optimizing, and Troubleshooting Power Related Parameters in Electronics Systems", Nov 18, 2014