PCB Design 007 QuietPower columns

Systematic Estimation of Worst-Case PDN Noise

Target Impedance and Rogue Waves

Istvan Novak, Oracle, November 2015

In the 'dark ages' of power distribution design, the typical advice was to use a bulk capacitor and one 0.1uF bypass capacitor for every power pin on the digital circuit. This was very unscientific, but served the industry reasonably well in low-density and low-speed circuits. As the designs got more demanding, the target impedance concept was developed [1]. Using a target impedance, designers had a metric and a design goal to guarantee that the voltage transients stay within specified limits. Strictly speaking, the target-impedance concept is valid only for flat self-impedance profiles, however, most of our practical designs do not have that. With non-flat impedance profiles, the noise is different; surprisingly and un-intuitively, keeping the same maximum impedance, the more we deviate from the flat impedance by pushing the impedance down in certain frequency ranges the higher the worst-case transient noise becomes. This raises the question how to do a systematic design and also gave rise to speculations about rogue waves [2]. But there is a systematic, fast and efficient way of calculating the worst-case noise for any arbitrary impedance profile.

The target impedance concept assumes that the power distribution network is hit by a series of current steps, each current step having a magnitude of ΔI and fastest transition time of t_{tr}. If up to the *BW* bandwidth of the excitation the PDN impedance is Z_{target}, the resulting voltage transients are within the ΔV limits.

$$BW = \frac{1}{\pi\, t_{tr}}$$

$$Z_{target} = \frac{\Delta V}{\Delta I}$$

The target-impedance concept and the above expressions assume a linear and time-invariant (LTI) PDN, moreover assume that the PDN impedance is flat, frequency independent, from DC up to the *BW* bandwidth of the excitation.

Interestingly, if the impedance profile is not flat, but still stays at or below the Z_{target} limit, the worst-case transient noise gets bigger. For some of the typical PDN impedance profiles this was shown in [3]. When the impedance profile is not flat and the worst-case transient noise is different from what we can expect from the target impedance formula, we need to determine what the excitation pattern that yields worst-case noise is and what its value is. Recently modified target impedance approaches have been proposed (see for instance [4]), or as [3] suggested, a conservative correction factor can be used based on

the degree of non-flatness of the impedance synthesis method. Using a conservative correction factor from the beginning makes it possible to follow a straightforward design process without the need of iterations.

For LTI PDNs with flat or any non-flat impedance, a process called *Reverse Pulse Technique* was published in 2002 [5]. Without the need of an optimization loop it provides a guaranteed way to determine the absolute worst-case transient noise and its corresponding excitation pattern, what we still may call rogue wave. To illustrate the power and usefulness of the process, we take the rogue-wave example circuit from [2] and calculate the worst-case noise with the *Reverse Pulse Technique*.

Figure 1 shows the schematics from [2], redrawn in a free circuit simulator [6]. Note that this particular simulator has the capability to represent a full RLC model of a single component, but for sake of clarity the schematic shown here explicitly calls out all parasitic elements and their own parasitics are set to zero. For example, component L_2, having an inductance value of 2nH, has no series resistance or parallel capacitance. The series resistance of L_2 is separately called out by R_2 with 2mOhm value.

We can run an AC simulation on this circuit to find out its impedance. For this purpose we run an AC sweep of the I_1 current source with a current magnitude of 1A. The V_2 voltage source with a voltage of zero is included only for convenience so that we can also plot the current going through our circuit. The voltage as a result of the 1A swept-frequency sine-wave excitation gives us the complex impedance. *Figure 2* shows the impedance magnitude and phase at node Z_{out} in the frequency range of 100Hz and 1GHz. The sweep was logarithmic with 100 frequency points per decade.

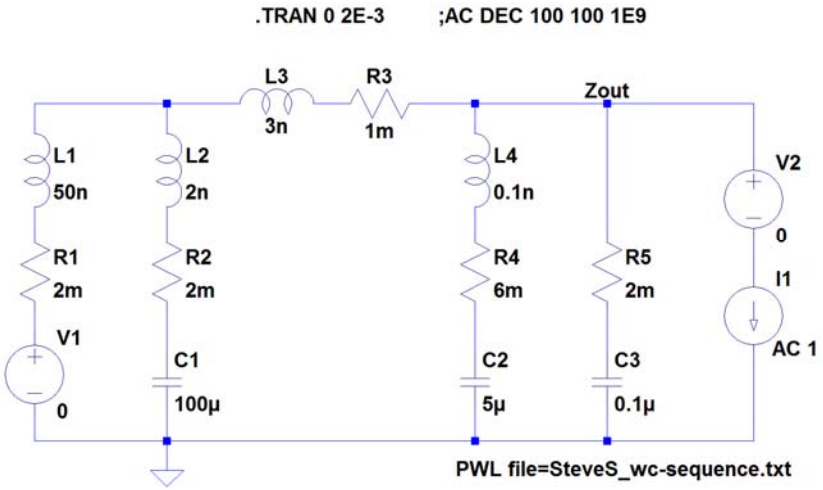

Figure 1: *Rouge-wave example circuit from [2].*

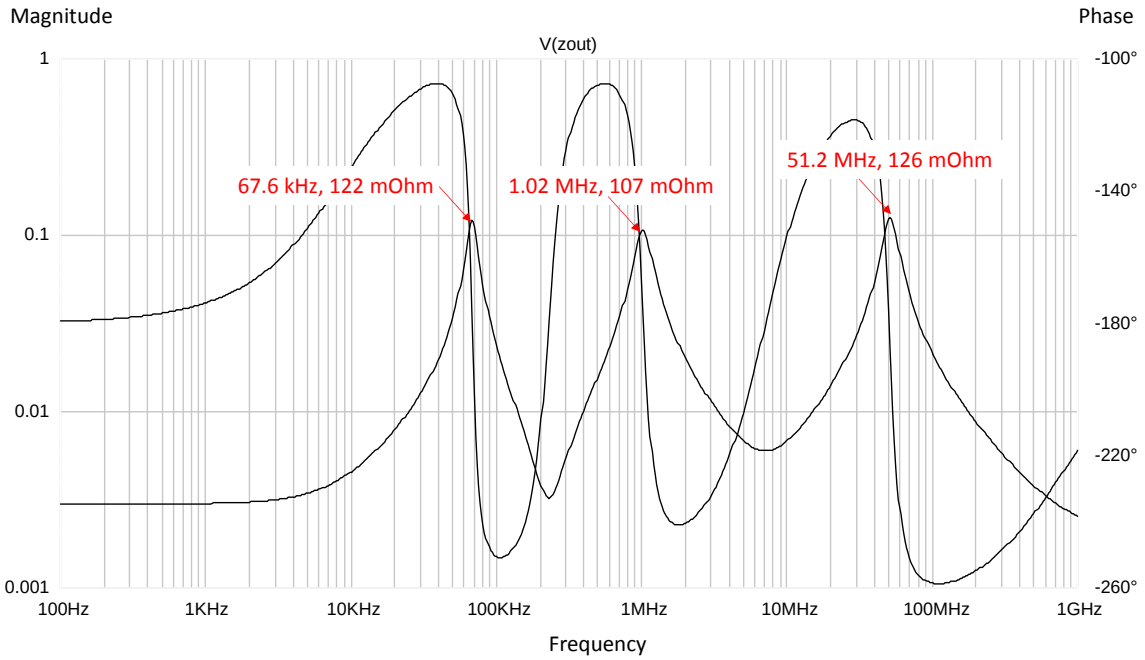

Figure 2: *Impedance magnitude and phase from the circuit shown in Figure 1. Note that both axes are logarithmic; in particular, the frequency scale is logarithmic to clearly show the resonance peaks separated by three orders of magnitude.*

The impedance profile shows three peaks with almost the same peak value, all slightly above 100 mOhm. The series loss values are very low, 1 – 6 mOhm, resulting in deep valleys in between the impedance peaks. The antiresonance frequencies spread across almost three decades of frequencies. While this impedance profile would be very rare in practice, and would most likely be the result of either careless design or lack of any systematic design whatsoever, we cannot rule out either the possibility that this could represent an actual circuit. Making use of the three distinct peaks separated by deep valleys, [2] uses a semi-heuristic approach to find what is called a rogue wave: it defines three repetitive bursts hitting the peak impedances one after the other, leaving the timing adjustment to an optimizer to find the biggest noise. The result is 750 mVpp for a series of 2A current step, which is equivalent to 375 mVpp/A. Compared to a perfectly flat impedance profile matching the largest peak, 126 mOhm, the optimization from [2] predicts a worst case of almost exactly three times of that value. The question is: is this really the worst case, or is it possible to find a different sequence of current steps that would produce an even bigger transient noise? We can turn to the *Reverse Pulse Technique* to get the answer.

The *Reverse Pulse Technique* starts with the *Step Response* of the circuit. Since the basic assumption is that the PDN is linear and time invariant (LTI), it does not matter whether we look at the response for a positive or negative going current step excitation; they are mirror images of each other. *Figure 3* shows the *Step Response* for a positive going current step. Without restricting generality, we assume that the DC voltage on the supply

rail is zero and therefore most of the transient response will be negative. Because of the LTI assumption, any DC voltage on the rail can be taken into account as a simple shift.

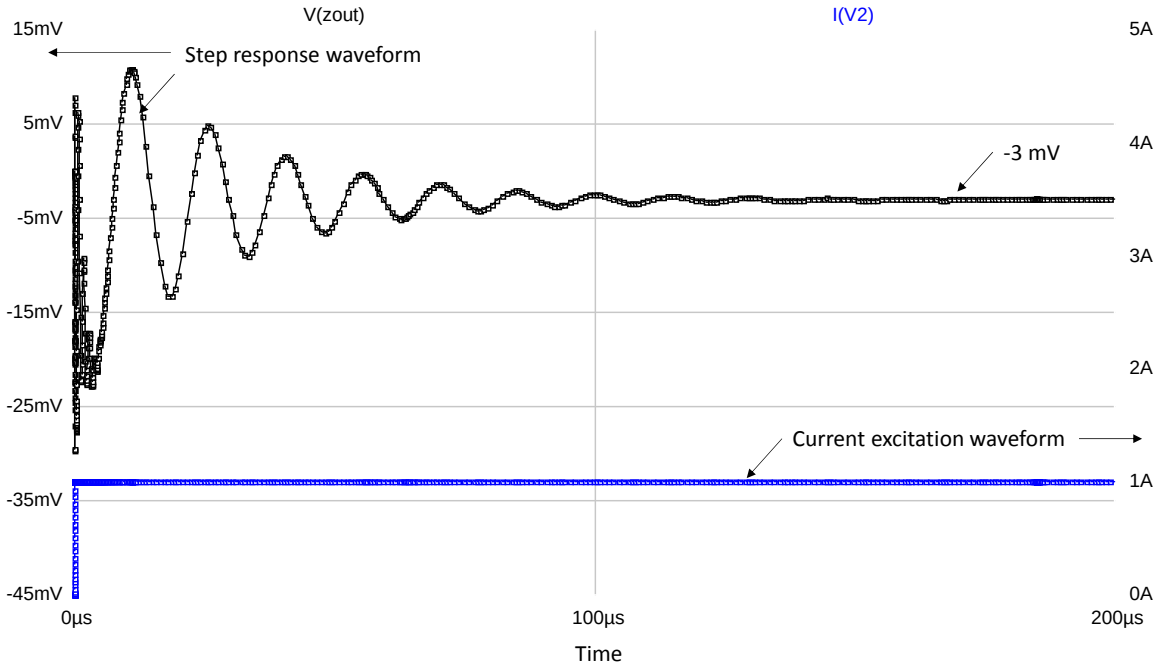

Figure 3: *Simulated Step Response of the circuit shown in Figure 1. Note that both axes are linear, as it would be shown on an oscilloscope or by a default simulation setup. Horizontal axis shows the full 0 to 0.2 ms time interval.*

The excitation current is a single current step, stepping up from zero to one ampere with 1 ns rise time. Note that the 1 ns rise time corresponds to about 300 MHz excitation bandwidth, where the impedance profile has a capacitive downslope and therefore the actual rise or fall time of the excitation is less critical.

The *Step Response* is shown up to 200 us, where it settles out to a -3 mV DC value. This is the result of the series equivalent of R_1 and R_3. The main signature we see on this scale is a damped sinusoid ringing with approximately 15 us period; this corresponds to the 67.6 kHz lowest-frequency peak in the impedance profile. We see more rapid changes near the left vertical axis, but we don't see any details. To see more details of the faster transients, we need to change the time scale. *Figure 4* shows the result. It is the same data, except now we show only the first ten microsecond time interval. On the right half of the plot the *Step Response* has a smooth rise; this is the beginning of the 67 kHz ringing. The left half of the response has a damped sinusoidal ringing with approximately 1 us period; this is originated from the 1.02 MHz impedance peak. We see some further fast transients near the left vertical axis, but on this horizontal scale we still cant see the details. We have to make another adjustment to the horizontal scale to see those details as well. On *Figure 5* we further zoom into the waveform and show only the first one microsecond interval. From 0.1 to 1 us we see a slow sine-wave in the response; this is the 1MHz damped sinusoid. Near the left vertical axis now we see

another damped sinusoidal waveform with approximately 20 ns period; this comes from the 51 MHz peak.

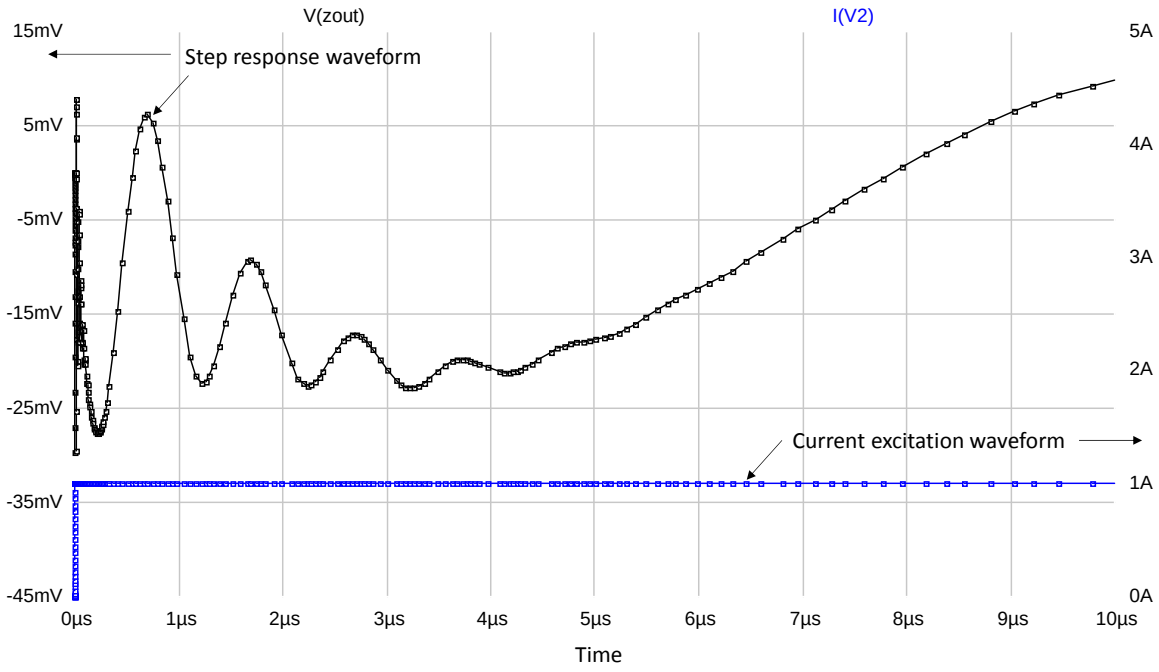

Figure 4: *Simulated Step Response of the circuit shown in Figure 1. Both axes are linear. The horizontal axis shows the first 0 to 10 us time interval.*

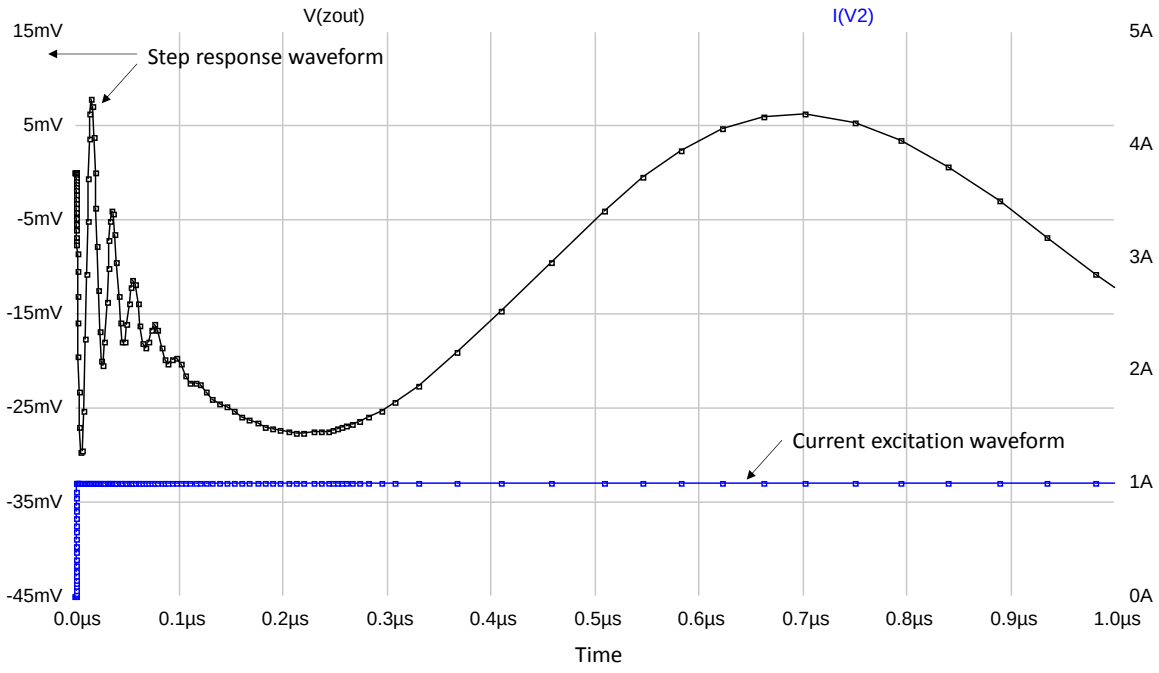

Figure 5: *Simulated Step Response of the circuit shown in Figure 1. Both axes are linear. The horizontal axis shows the first 0 to 1 us time interval.*

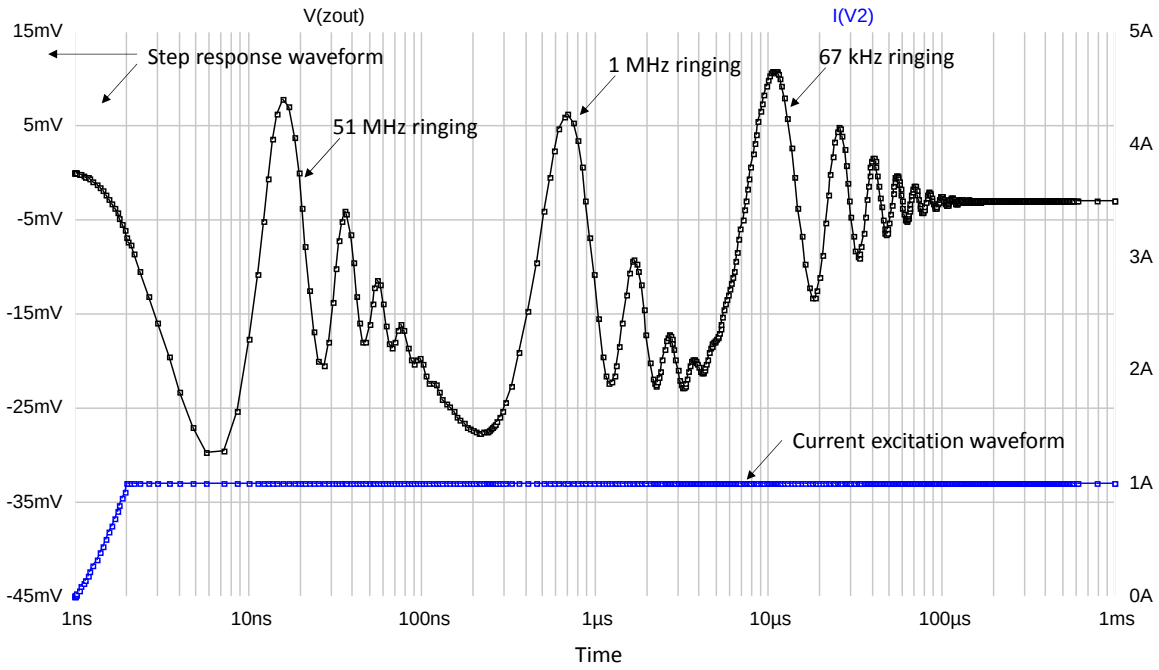

Figure 6: *Simulated Step Response of the circuit shown in Figure 1. Vertical axis is linear, the horizontal axis is logarithmic.*

To see all signatures on the same plot, we need to switch to logarithmic horizontal axis, as shown in *Figure 6*. The logarithmic time axis, just like the logarithmic frequency axis on the impedance plot, allows us to see very different signatures on the same plot. We now clearly see side by side all three damped sinusoid responses.

With the *Step Response* data in *Figures 3* through *6* we can continue the process of the *Reverse Pulse Technique*. (Note that *Figures 3* through *6* show the same exact data only in different forms) Next we have to identify the steady state and the peaks and valleys in the *Step Response*. We have to do it in reverse order, starting with the right-most first extremum (peak or valley) and step through the peaks and valleys one by one from right to left until we reach the excitation time instance. *Figure 7* shows the time stamps and voltage values of the peaks and valleys identified in the *Step Response*. Note that in simulated waveforms, like in this case, identifying the peaks and valleys automatically is relatively easy; it would become more difficult when we need to process *Step Response* waveforms obtained by measurements. The measurement noise makes the peak/valley identification a little trickier.

With the data points in *Figure 7* we can continue in two different ways. If we do not need to identify the pattern of the rogue wave excitation and we need only the worst-case transient noise magnitude, we just need to sum up the peaks and valleys and take the difference. The sum of the peaks is -78 mV; the sum of the valleys is -275 mV. The difference is -197 mV. The -197 mV value is the absolute worst-case one-sided noise when an arbitrary sequence of 1A current steps hits the PDN. The worst-case two-sided

transient noise is twice of this value minus the DC steady-state value (-3 mV in this case). These numbers give us 391 mVpp worst-case transient noise. The other possible way of continuing with the data points from *Figure 7* is to determine the time-domain sequence of excitation edges creating the worst-case noise (the rogue wave) and to actually simulate the time-domain noise.

Time [ns]	Peak [mV]	Valley [mV]
137012.9		-3.1
129952.8	-2.8	
122328.4		-3.2
116025.0	-2.7	
107579.3		-3.4
99748.4	-2.5	
92487.5		-3.7
85109.5	-2.1	
77730.4		-4.2
69926.0	-1.4	
62905.3		-5.1
55320.8	-0.3	
47920.9		-6.5
40274.7	1.6	
33340.8		-9.0
25591.5	4.9	
18351.7		-13.3
10894.3	10.8	
4172.1		-21.2
3724.9	-19.8	
3226.7		-22.8
2691.4	-17.2	
2262.0		-22.5
1697.3	-9.3	
1226.4		-22.2
701.0	6.2	
217.3		-27.6
97.5	-19.7	
89.7		-20.3
76.6	-16.2	
67.8		-18.6
56.2	-11.4	
47.6		-18.0
36.0	-4.1	
27.0		-20.4
15.9	7.9	
6.3		-29.6

Figure 7: *Peak and valley time stamps and voltages identified in the Step Response of the circuit shown in Figure 1.*

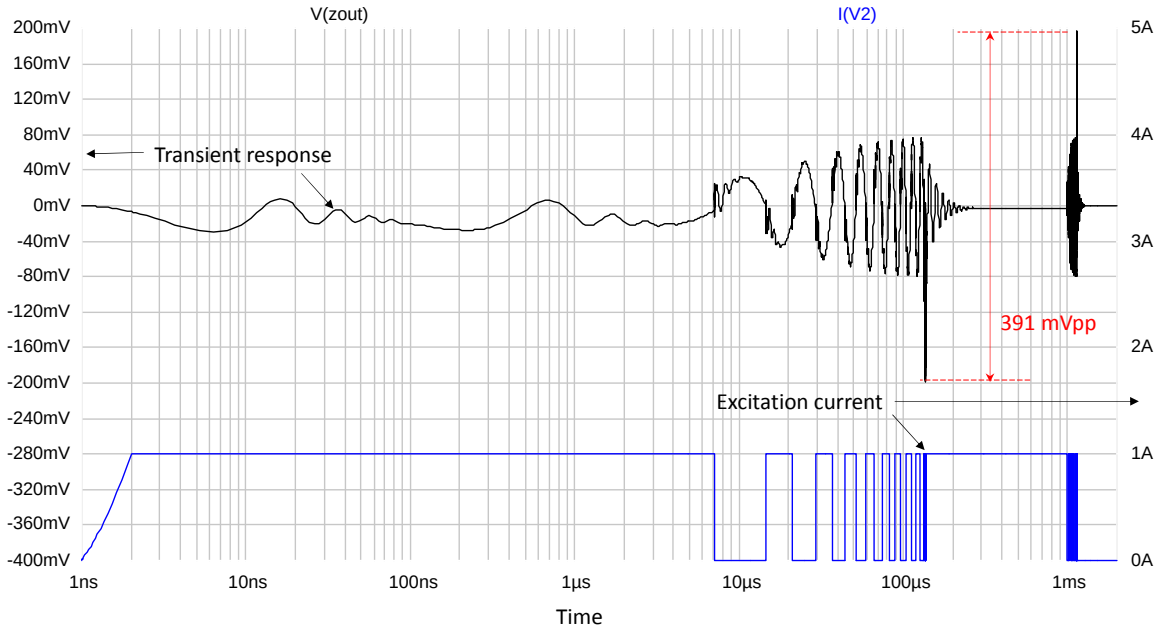

Figure 8: Worst-case response simulated with an excitation sequence calculated from the Reverse Pulse Technique.

With the timing sequence from *Figure 7*, *Figure 8* shows the simulated waveforms on logarithmic horizontal scale. The blue waveform on the bottom is the excitation waveform, the black waveform on the top is the transient response. The peak-to-peak transient value is 391 mVpp, exactly matching the value that we calculated just from the peaks and valleys of the *Step Response*. Note that to achieve the worst-case transient noise, we used 37 current steps and their spacing does not exactly follow the three resonance frequencies. This straightforward process yields the worst-case noise very fast, without the need of an optimization loop and it guarantees to provide the worst-case noise. In this particular example the true worst-case noise is 391 mVpp/A as opposed to the 375 mVpp/A predicted by the rogue-wave optimization from [2].

We can also look at the transient noise by another popular test method: using a repetitive stream of current steps with 1A magnitude and tune the repetition frequency (and possibly also the duty cycle) until we observe the maximum noise. We just change the definition of the I_1 current source to a stepped-frequency square-wave and run the simulations again. As we change the repetition frequency, we find that we get the maximum noise magnitude when the repetition frequency matches one of the peak frequencies. *Figure 9* shows the result when we move the frequency around the middle peak in five values: 300 kHz, 900 kHz, 1020 kHz, 1100 kHz and 5000 kHz.

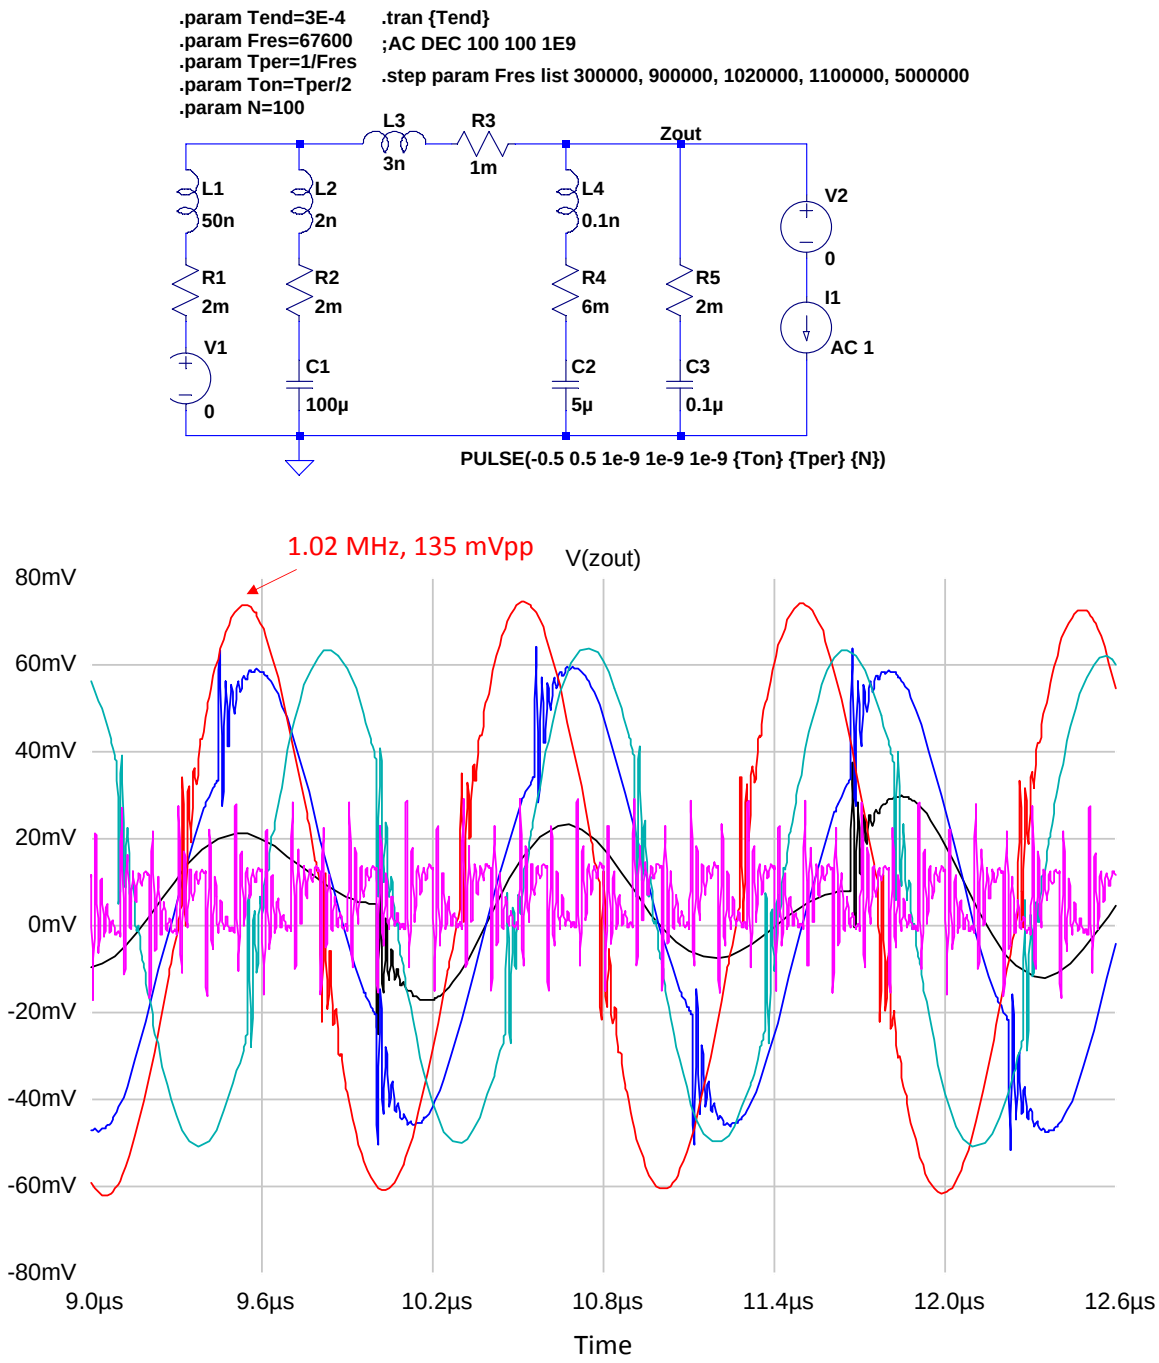

***Figure 9**: Schematics with the excitation definitions on the top, response waveforms on the bottom. The square-wave excitations frequencies map out the middle resonance. The biggest response comes from the 1020 kHz square wave, matching the middle resonance frequency in the impedance profile.*

Figure 10 shows the schematics and waveforms for the three repetition frequencies exactly matching the three peak frequencies. We see that in this case the transient noise is higher, actually $4/\pi$ times higher, than the product of the impedance peak magnitude

and current magnitude. This is because the high-Q peak picks out the fundamental harmonic and greatly attenuates the harmonics. In the Fourier series of a square wave with 50% duty cycle, the fundamental-frequency sine wave has a $4/\pi$ times higher magnitude than that of the square wave.

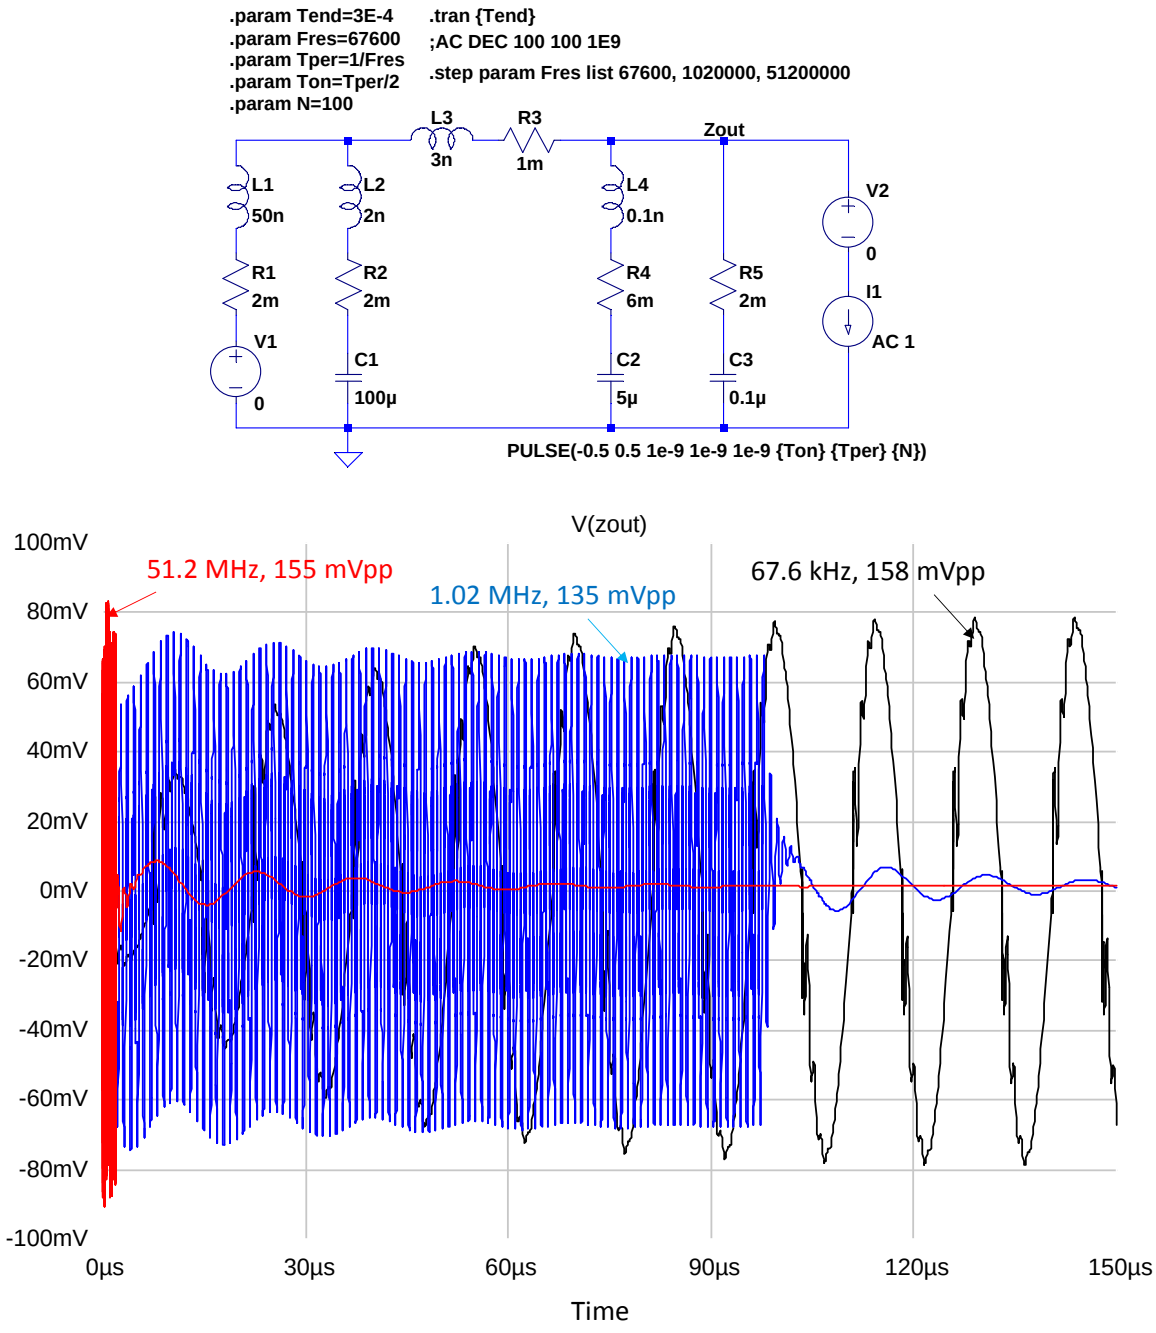

***Figure 10**: Schematics with the excitation definitions on the top, response waveforms on the bottom. The square-wave excitations frequencies exactly match the three resonances.*

Finally the table in *Figure 11* summarizes the characteristic noise signature numbers we obtained. Note that the *Reverse Pulse Technique* yielded the highest noise, and it is proven to be the absolute worst case. It is also true that the rogue-wave optimization could provide the same (correct) answer, however, in a multi-resonance case like this example, without operator guidance it could take a lot of computing resources and eventually it may not converge.

Worst-case transient noise estimate calculated from:	Peak-to-peak mV	% error
Step Response peak value and steady state	56.2	-85.7
Swept-frequency periodic pulse stream	158	-59.6
Rouge wave optimization	375	-4.1
Reverse Pulse Technique	391	0

Figure 11: *Worst-case transient noise estimates of the circuit in Figure 1 based on different calculation methods. All responses assume a sequence of 1A current steps.*

The first entry in the table is calculated as twice the peak deviation of the *Step Response* (from the last row in *Figure 7*) minus the steady state response. This is the peak-to-peak noise as a result of a single rising edge followed by a single falling edge with a large time separation in between. This estimate is 85.7% smaller than the true maximum. The second entry equals the biggest peak in *Figure 2* (126 mOhm at 51.2 MHz) multiplied by $4/\pi$. This is the result of sweeping a periodical current pulse stream to find the maximum noise deviation. This estimate is 59.6% lower than the true worst case. The rouge-wave optimized value is taken from [2]. It is 4.1% lower than the true worst case.

In this particular example, when the impedance profile has multiple, almost equal peaks, the difference is dramatic. Estimating the worst-case noise just from the peak deviation or from swept-frequency periodic excitation hugely under-estimates the worst-case noise. The rogue-wave optimization, in theory, should be able to find the true worst case, but at a price of significant run time and potential convergence failures. With a more flat impedance profile, with fewer peaks and smaller peak-valley ratio, the errors in all of the approximations would be lower. Eventually for a perfectly flat impedance profile all four calculation methods would provide the same result.

If you want to learn more about the subject, follow [7] and [8] at DesignCon 2016.

References:

[1] Larry D. Smith, Raymond E. Anderson, Douglas W. Forehand, Thomas J. Pelc, and Tanmoy Roy, ‘‘Power distribution system design methodology and capacitor selection for modern CMOS technology’’, IEEE Transactions on Advanced Packaging, vol. 22, no. 3, pp. 284-291, Aug.1999.

[2] Steve Sandler, “Target Impedance Limitations and Rogue Wave Assessments on PDN Performance,” paper 11-FR2 at DesignCon 2015, January 27 – 30, 2015, Santa Clara, CA.
[3] Istvan Novak, “Comparison of Power Distribution Network Design Methods: Bypass Capacitor Selection Based on Time Domain and Frequency Domain Performances,” DesignCon 2006, February 6 – 9, 2006, Santa Clara, CA.
[4] Jingook Kim, Songping Wu, Hanfeng Wang, Yuzou Takita, Hayato Takeuchi, Kenji Araki, Gang Feng, and, Jun Fan, “Improved Target Impedance and IC Transient Current Measurement for Power Distribution Network Design,” Proceedings of the IEEE EMC Symposium, July 25 – 30, 2010, Fort Lauderdale, FL.
[5] Drabkin, et al, “Aperiodic Resonant Excitation of Microprocessor power Distribution Systems and the Reverse Pulse Technique,” Proceedings of EPEP 2002, p. 175.
[6] http://www.linear.com/designtools/software/#LTspice
[7] Target Impedance and Rogue Waves, panel discussion at DesignCon 2016, January 19 – 21, 2016, Santa Clara, CA.
[8] Jae Young Choi, Ethan Koether, Istvan Novak, “Electrical and Thermal Consequences of Non-Flat Impedance Profiles,” DesignCon 2016, January 19 – 21, 2016, Santa Clara, CA.