DesignCon East 2005

Bypass Filter Design Considerations for Modern Digital Systems, A Comparative Evaluation of the Big "V", Multi-pole, and Many Pole Bypass Strategies"

Steve Weir, Teraspeed Consulting Group
steve@teraspeed.com
(401) 284-1827

Abstract

Ceramic capacitor bypass filter networks may be constructed by any of three well-known methods popularly differentiated by the number of capacitor values used: one, one per capacitor per decade, or multiple capacitors per decade. Support of a given method seems to be nearly religious in fervor. We develop a general model for all three methods that provides insights into the strengths and weaknesses of each method from all: power delivery, signal integrity, EMC and manufacturing perspectives. Our goal is to reduce religious debate to practical choices based on application criteria. Finally, we demonstrate efficient network synthesis for each technique.

Author(s) Biography

Steve Weir, Consultant, Teraspeed Consulting Group

Steve is an independent consultant with over 20 years plus industry experience with a broad range of expertise. Steve holds 17 US patents, and has architected a number of TDM and packet based switching products, consults on patents and is a frequent contributor to the SI-list signal integrity reflector.

What Does the Bypass Network Do?

A bypass filter network applies a shunt across the power rails of sufficiently low impedance to maintain rail voltage in the presence of switching currents. In a modern system, we are concerned with switching currents that range:

1. One time surge as the power rails initial charge. Large spikes can sustain for a few uS to a millisecond or more as the rails transition through a range that biases CMOS FETs in their linear regions. Duration is independent of rise-time, which can be in the low ns.

2. Repetitive surges as power-managed devices enable or disable large functional blocks.

3. Pulsating core currents associated with large state machine or memory block operations.

4. Pulsating I/O currents associated with signaling.

5. Where Vdd planes act as signal return image planes, bridges signal switching currents.

The ceramic capacitors in an application provide the necessary shunt impedance from a cut-off point where the voltage regulator module impedance rises above the system target, and either:

1. For systems that do not use bypass capacitors to bridge signal return currents between planes- a point somewhat above the package low-pass filter cut-off of the ICs.
2. For systems that do use bypass capacitors to bridge signal return currents between planes- the lesser of:
    a. Signal knee frequency,
    b. Parallel resonant frequency, PRF, where plane cavity net capacitive reactance crosses the discrete capacitor network net inductive reactance.

IC package power low-pass cut-off rarely extends above 100MHz and is fundamentally limited by package size. However signal currents routinely exceed 1GHz today and are headed up. It is reasonable to question the prudence of relying upon the bypass capacitors and plane cavities to convey image return currents between multiple reference planes.

Despite the low-pass nature of IC packages, massive modern IC currents still translate to very significant package ⇔ PCB currents at high frequency. These currents can become real EMC headaches when coupled onto networks, and / or structures that resonate or just provide a high impedance (read efficient antenna) near the excitation frequency.

MLCC Bypass Capacitor Basics

Capacitor impedance magnitude follows a familiar “V” shape. Impedance falls from a theoretical value of infinity at DC to a minimum representative of the ESR at device mounted self-resonant frequency where the capacitive reactance and inductive reactance are equal. At higher frequencies the inductive reactance dominates.

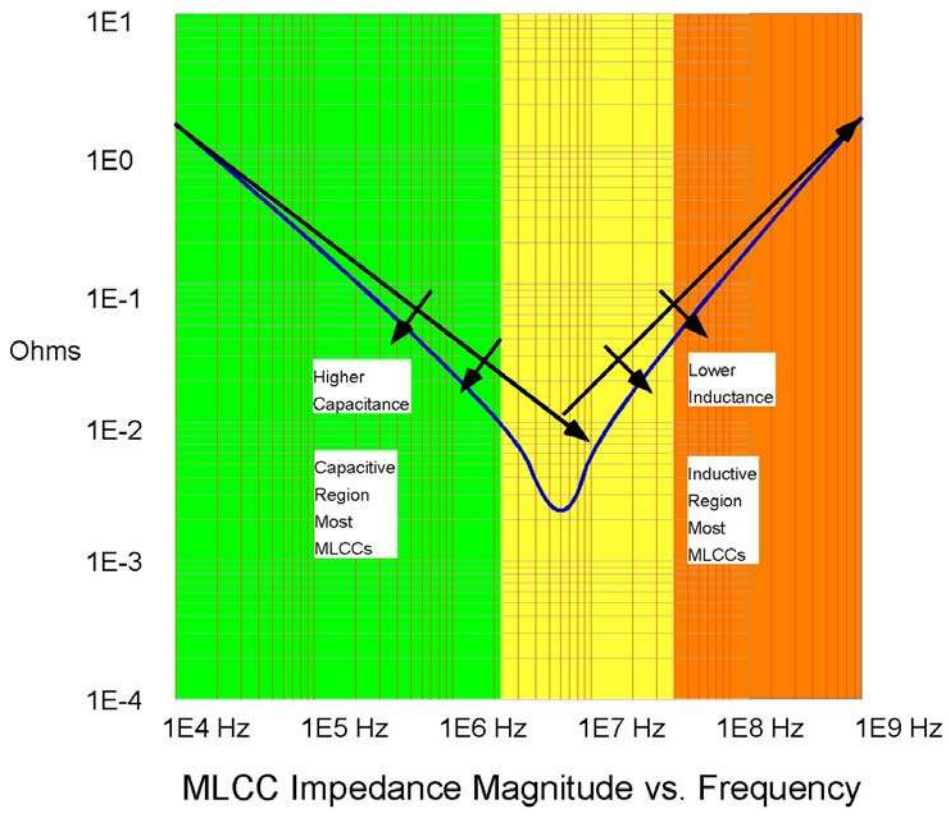

Figure 1, MLCC Impedance vs. Frequency (Lumped Model)

MLCC capacitors typically span capacitance ranges of 1000:1 or more in a given case size.

Typically available parts in 6.3V ratings:

- 0402 100pF – 470nF
- 0603 180pF – 2.2uF
- 0805 180pF – 10uF

Within any case size and dielectric composition from a given manufacturer, device ESR follows capacitance as:

Equation 1 $\mathbf{ESR = K1 * C^{K2}}$ **where -0.5 < K2 < -0.3**

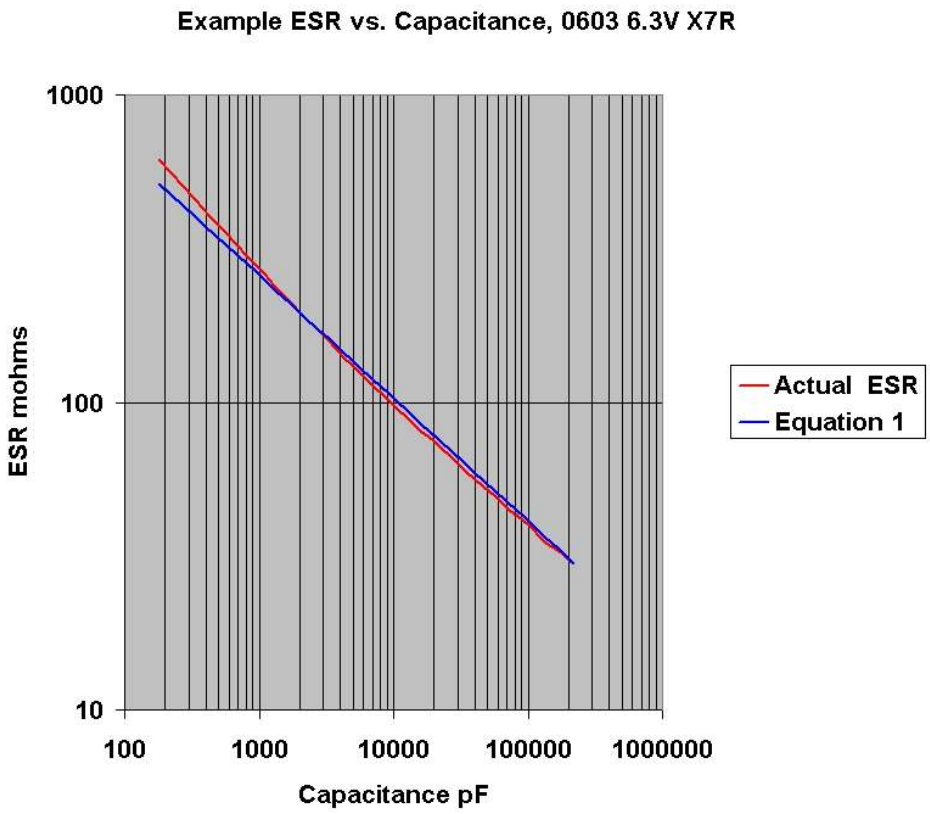

Figure 2, ESR vs. Capacitance, AVX 0603 X7R Values from SpiCap 3.0™[1] Shown

Parallel Resonance Between Capacitor Networks

Where the impedance magnitudes of two networks intersect, the phase angle difference between the two networks determines the severity of impedance peaking. This is of concern for MLCC bypass networks in three regions:

- Transition from the VRM / bulk capacitor network to MLCC network
- Transition from the MLCC network to board cavity capacitance
- With multi-pole MLCC networks: transitions between each capacitor value network

[1] SpiCap 3.0 ©2003 AVX Corporation

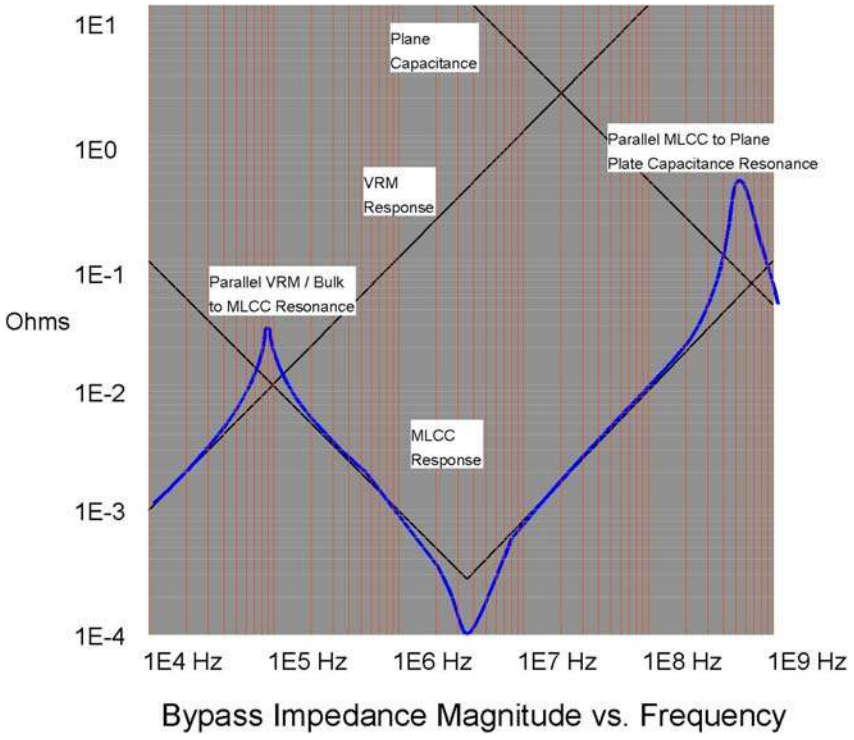

Figure 3, Bypass Network Impedance vs. Frequency

Given a series of capacitors with widely spaced pole pairs, the impedance magnitude crossover points occur at near 180 degrees phase difference, resulting in the following relationships:

Equation 2 $\omega_{PRF} \approx 1/(L_{LF\_NET} * C_{HF\_NET})^{0.5}$

Equation 3 $|Z_{LPRF}| \approx 1/(L_{LF\_NET} * C_{HF\_NET})^{0.5} * L_{LF\_NET} = L_{LF\_NET}^{0.5} * C_{HF\_NET}^{-0.5} = |Z_{CHAR}|$

Equation 4 $Q \approx |Z_{LPRF}| / (ESR_{HF\_NET} + ESR_{LF\_NET}) = |Z_{CHAR}| / (ESR_{HF\_NET} + ESR_{LF\_NET})$

Equation 5 $|Z_{PEAK}| \approx |Z_{CHAR}| * Q \approx |Z_{CHAR}^2|/(ESR_{HF\_NET} + ESR_{LF\_NET})$

If we set the capacitance of a given network as the sum of a group of capacitors, ie $C_{HF\_NET} = C_{HF\_CAP} * N$ then:

Equation 6 $ESR_{HF\_NET} \approx ESR_{HF\_CAP} / N$

Equation 7 $|Z_{CHAR}| \approx L_{LF\_NET}^{0.5} * C_{HF\_CAP}^{-0.5} * N^{-0.5}$

Equation 8 $Q \approx L_{LF\_NET}^{0.5} * C_{HF\_CAP}^{-0.5} * N^{0.5} / (ESR_{HF\_CAP} + N * ESR_{LF\_NET})$

Equation 9 $|Z_{PEAK}| \approx |Z_{CHAR}| * Q \approx L_{LF\_NET} / (C_{HF\_CAP} * (ESR_{HF\_CAP} + N * ESR_{LF\_NET}))$

Equation 10 $|Z_{PEAK}| \approx L_{LF\_NET} / (C_{HF\_CAP}^{(1+K2)} * K1)$

Since K1, and K2 are constant values that depend on the capacitor manufacturing process, and L_{LF_NET} is set by the lower frequency network, Equation 10 implies that in order to contain peak impedance, we should seek the largest capacitance value in a given package / geometry combination. Figure 4 illustrates this principle as it applies to a single capacitor. Figure 5, illustrates that scaling capacitor quantity provides no improvement to $|Z_{PEAK}|$ when the result fails to improve phase margin.

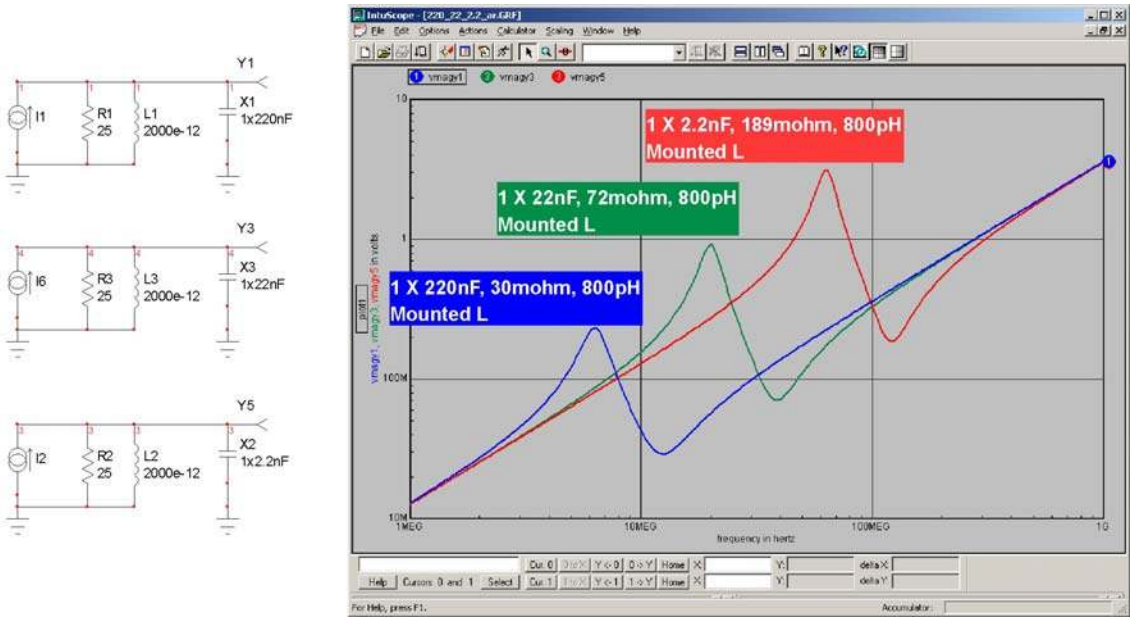

Figure 4, Example Parallel Resonance

Equation 10 also reinforces the concept that our network impedance is driven by: inductance, inductance and inductance of all the elements in the network, not just the high frequency capacitors. At equal cost, capacitors with lower mounted inductance afford benefits across the frequency spectrum.

Figure 5 illustrates that when the impedance magnitudes intersect with a phase difference close to 180 degrees, that peak impedance remains essentially constant, while Q increases markedly with increasing capacitor count. Figure 6, illustrates that we can bring parallel resonant peaking under control only by establishing the magnitude crossings with substantial phase margin versus 180 degrees.

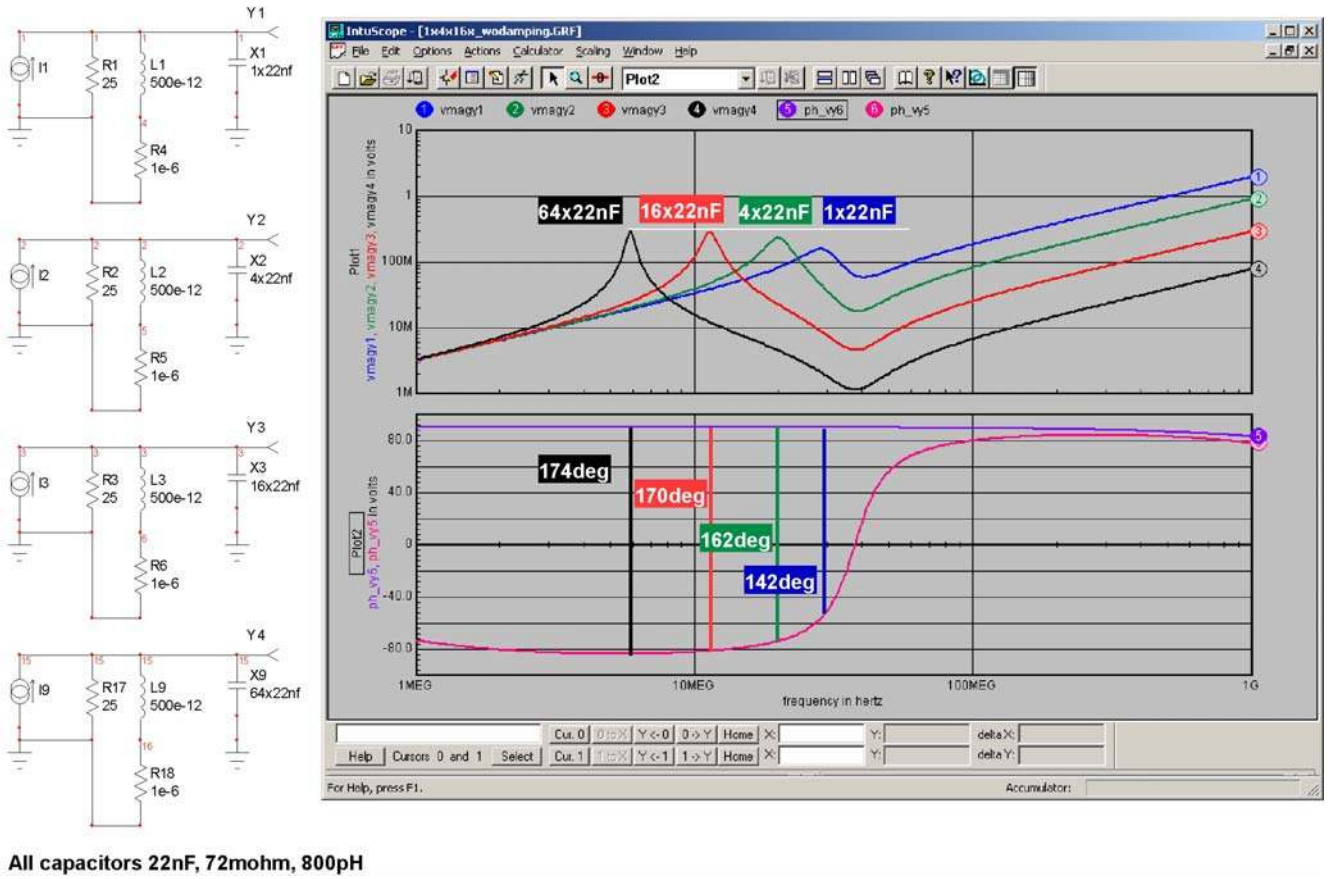

Figure 5. Example Peak |Z| versus Capacitor Count, Low Phase Margin

In order to manage peak impedance, we need to provide phase margin at the impedance magnitude intersections. Figure 6 illustrates the tremendous difference phase margin makes:

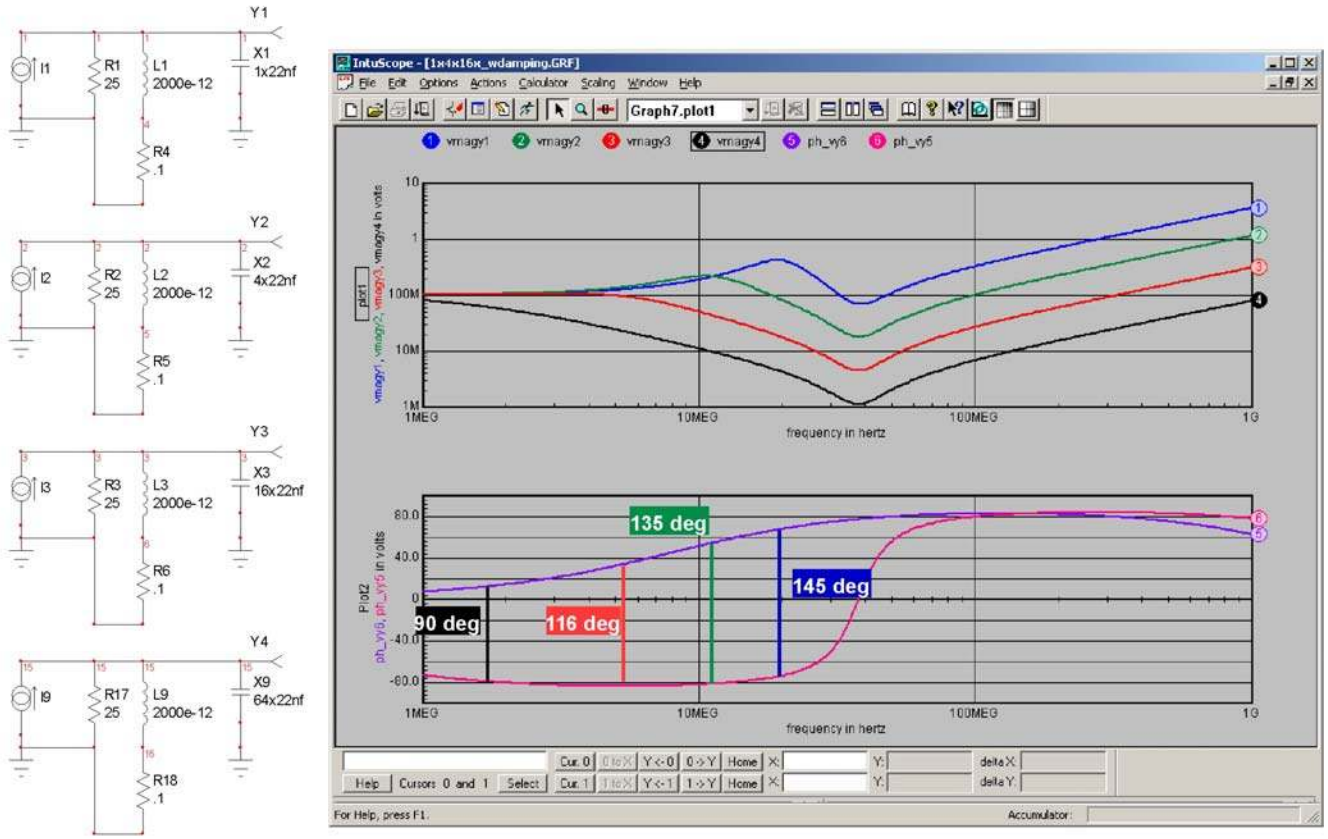

Figure 6, Example Peak |Z| versus Capacitor Count, High to Low Phase Margin

In this example, the 0.1 ohm resistors limits the phase shift of the L/R network until the 10MHz region. So long as we cross impedance magnitude with at least 45 degrees phase margin, impedance peaking will be limited to 25% or less of |Z| or either network alone at the transition.

Bypass Strategies-Three Methods, Three Faiths?

Today, three methods of bypass strategy may be considered popular:

- Finely spaced multipole, (Sun Microsystems, Larry Smith et-al, flat response)
- Coarsely spaced multipole, (“Capacitors by the decade”)
- Single value ceramic capacitor, (big “V”)

Multi Pole Filters, F=K*X^N Versus Single Value Big “V”

Beginning in the late 1990s, researchers at Sun Microsystems published a series of papers[1] on a novel method of bypass network design that concatenates many tightly spaced filters. The method emphasizes derivation of high frequency network performance dictated more by ESR, than by ESL.

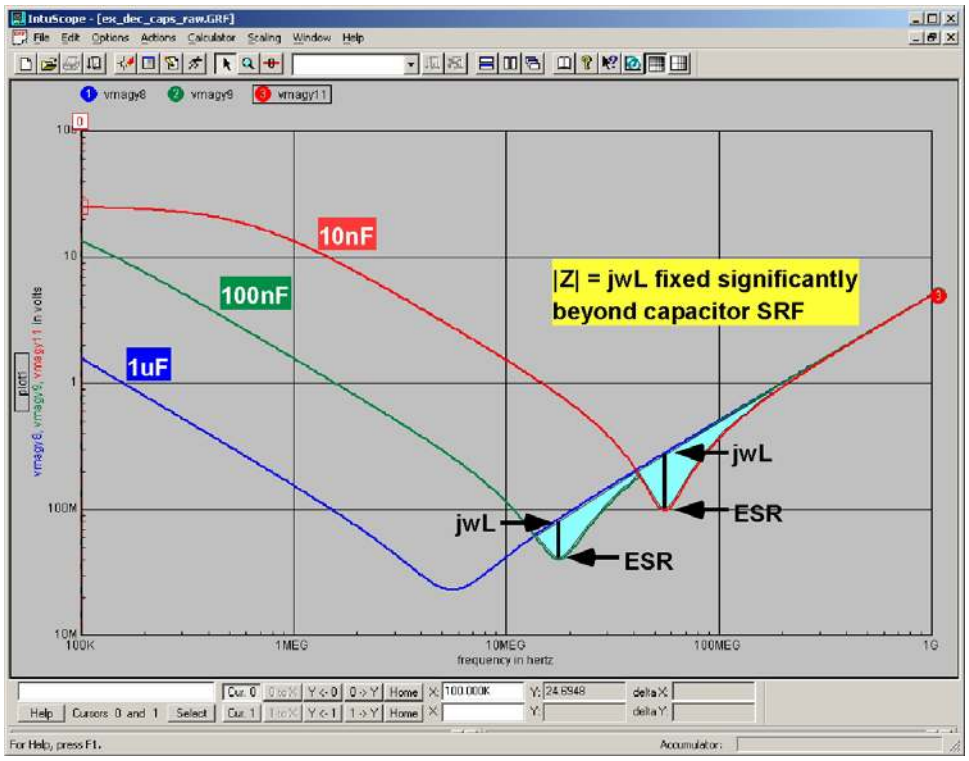

Figure 7, ESR Impedance[2] at SRF Versus Mounted Inductance

As can be seen from Figure 7, MLCC capacitors, particularly in smaller capacitance values exhibit high Qs: the ESR at the SRF for each capacitor is much lower than jωL. Sun researchers reasoned: given a particular impedance target and bandwidth, one should be able to capitalize on the high Q of these devices and develop networks dominated by the ESR value rather than jωL. The resulting networks should have desirable characteristics when compared to other methods:

- Reduced component count
- Higher ESR, which would reduce peaking at parallel resonance with PCB plane cavities
- Faster transient response

The method takes good advantage of low relative phase angles that result with closely spaced SRFs to limit the parallel resonant peaking between networks. If one wished to go to extremes the entire E12 capacitor series could be used.

However some balance between benefits of large phase margins and manufacturing complexity is called for. At present, the method is often advocated with three capacitor values in each decade. Ie: 2.2nF, 4.7nF, 10nF, 22nF,[6] etc.

In order to maintain a fairly flat impedance floor, the number of capacitors required in each successively smaller capacitance value increases commensurate to ESR, plus[1]. Given a network primarily composed of capacitors all in the same package, and therefore all with the same mounted ESL, the SRF of each network follows the form:

Equation 11 $\omega_{SRF(N+1)} = \omega_{SRF(N)} \left(C_N / C_{N+1} \right)^{0.5}$

[2] Includes parallel 50 ohm source, and 50 ohm receiver

Equation 12 $X = (C_N / C_{N+1})^{0.5} = 10^{1/6}$

Ie, we have a series of networks each with a self-resonant frequency of $F = K*X^N$. For the case of three capacitors per decade, X = 1.47.

Assuming for a moment that K2 from Equation 1 exactly equals –0.5, then we would further have:

Equation 13 $ESR_{(N+1)} = ESR_{(N)}*X$

Equation 14 $j\omega L_{SRF(N+1)} = j\omega L_{SRF(N)}*X$

With the immediately observable result that $Q_{(N)}$: $j\omega L_{SRF(N)} / ESR_{(N)}$ is a constant.
The number of capacitors required to hold a constant ESR, and a constant $|Z_{CHAR}|$ in each successive value bin then reduces to:

Equation 15 $COUNT_{(N+1)} = ESR_{(N+1)} / ESR_{(N)} * COUNT_{(N)} = X * COUNT_{(N)}$

Given a capacitor count that is fixed by frequency, for a given capacitor mounted inductance, case size and chemistry, filter bandwidth simply extends by adding more high frequency capacitor values in quantities proportionate to the resulting SRFs. The only remaining unknown is how many capacitors to seed into the lowest frequency bin.

The seed quantity problem on paper is dictated by the extent of loading interaction between the parallel networks. Due to the constant Q, we can show that for K2 = 0.5, a simple constant is sufficient to scale ESR.

This math works independently of the value of X. In the case of the fine-grained networks advocates, X = 1.47, whereas for capacitors on decade values, X = 3.16.

Sample Networks

An illustration of a hypothetical one dimensional network using fractional capacitor counts, K2 = 0.5, and ESR_{220nF} = 30 mohms:

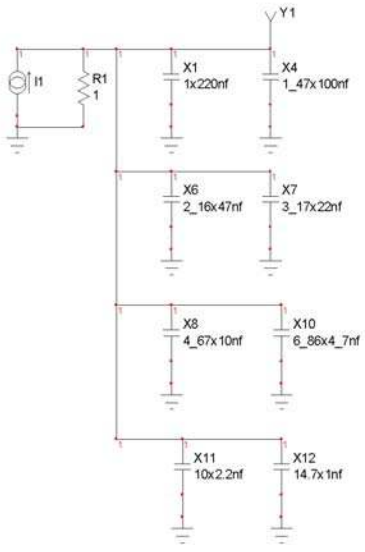

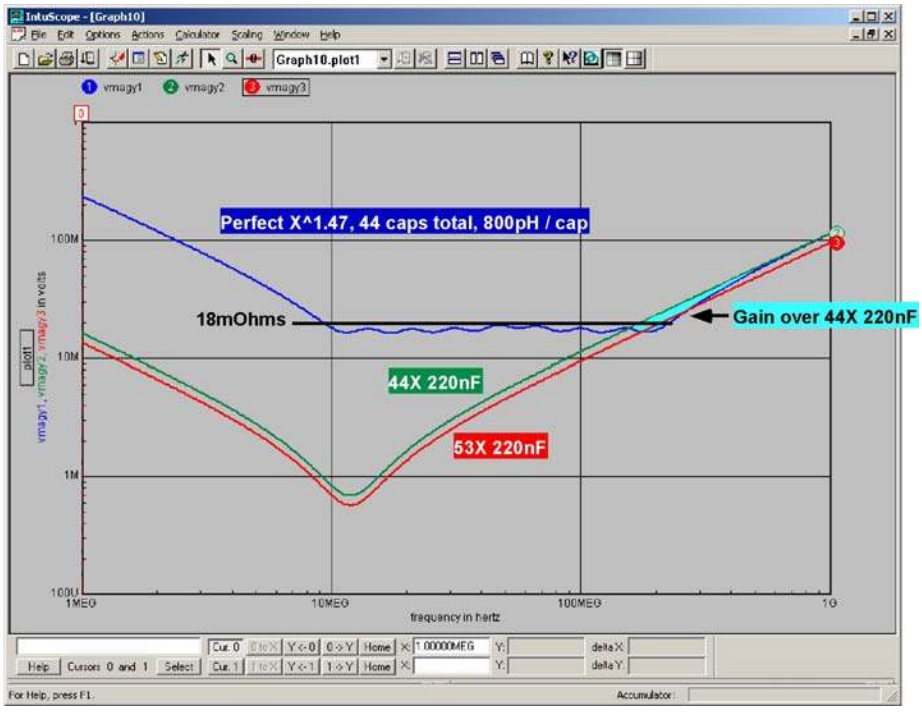

Figure 8, Perfect 1nF, 2.2nF, 4.7nF ... 220nF Network vs. 220nF big "V"

Table 1, summarizes the network.

Capacitor	ESR	Qty
220nF	30m	1
100nF	44.1m	1.47
47nF	64.8m	2.16
22nF	95.3m	3.18
10nF	140m	4.67
4.7nF	206m	6.86
2.2nF	300m	10.0
1nF	441m	14.7
Net	3.75m	44.04

Table 1, Example Perfect 1.47^N Network

We see that the impedance floor peaks occur at 18mohms: 60% of the ESR of each parallel capacitor network. The entire parallel ESR is 3.75mohms. From 160MHz to 500MHz, this network model delivers lower impedance than a network composed of 44 220nF capacitors only. To reach parity at the higher frequencies an additional nine capacitors, or 20% more are needed, 53 total using only 220nF capacitors. In each case, the single value networks exhibit ESRs only about one-fourth that of the complex network. Referring to Equation 9:

$$|Z_{PEAK}| \approx |Z_{CHAR}|*Q \approx L_{LF\_NET} / (C_{HF\_CAP} * (ESR_{HF\_CAP} + N*ESR_{LF\_NET}))$$

The higher ESR of this network versus either of the two big "V" networks results in a substantially lower impedance peak at the impedance crossover to the plane cavities for any crossover substantially above the big "V" network SRF. This is perhaps the greatest strength of the fine multi-pole approach.

This property of higher ESR, is potentially of significant importance in systems where the bypass network bridges signal return currents between image planes of both power rails, or where coherent signals such as clocks align close to the parallel resonant frequency. Figure 9 illustrates the comparison when we include plane cavities.

If however, all three conditions are met:

- IC power low-pass cut-off occurs well before the parallel resonant frequency, AND
- Signals do not excite near the resonant frequency sufficiently to present an EMC issue, AND
- The bypass capacitor network does not comprise a substantial portion of signal return paths

Under these combined conditions, the higher impedance peaks at the ceramic capacitor to plane cavity transition of big "V" implementations is of no consequence.

What about capacitors on decade spacing? If we synthesize a network to most closely match the performance of the 1.5^N using 3.16^N we get:

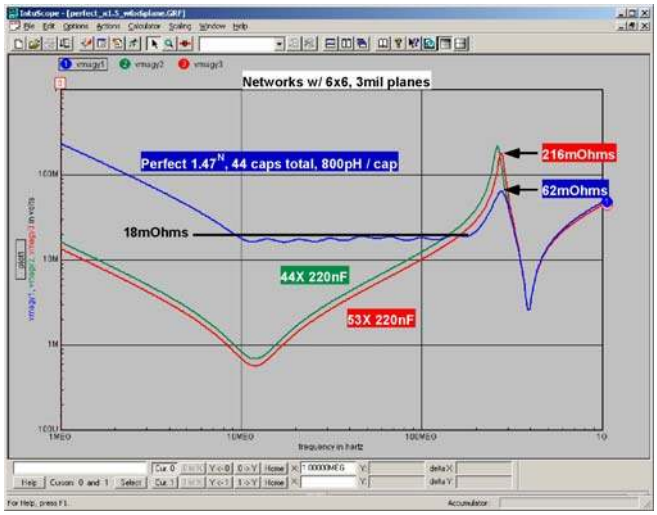

Figure 9, Perfect 1.47^N Network vs. Big "V", Including Planes

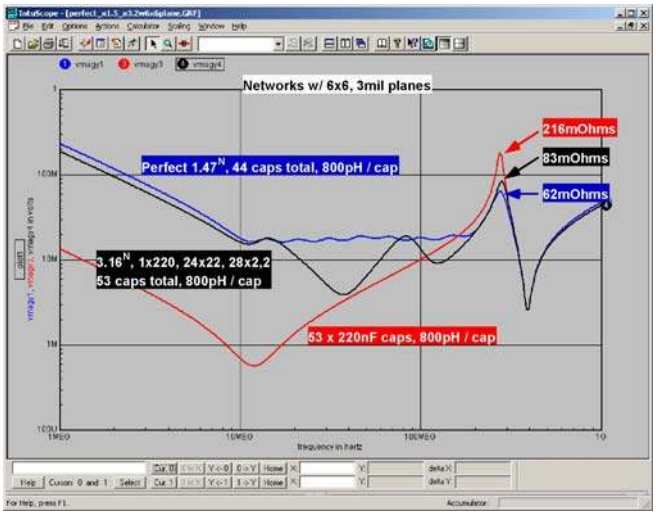

Figure 10, Comparison w/ Decade Spaced Capacitors

In order to match the bandwidth of the one-dimensional many pole filter model, a total of 53 capacitors, the same as the big "V" are required. What we buy for the extra complexity over the big "V", and extra component count over the many pole case, is the best impedance from 110 to 160MHz, slightly better impedance than the many pole case at the low end, and much better impedance at the parallel resonance with the plane cavity than the big "V" implementations. However, impedance is markedly higher than the big "V" over a wide range of frequencies, including the low frequency limit where we need to transition from the VRM / bulk capacitors.

The World in 2D

The foregoing one-dimensional models neglect spatial distribution of the capacitors, as well as plane transmission line effects. Both are very significant, and at high frequencies dominate power delivery performance. The component count advantage of fine multi-pole often fades when we account for spatial distribution of the capacitors and transfer impedance to the BGA balls. The effect is particularly acute for IC's with center power slugs, and / or relatively thick power plane cavities in the IC package.

Here we model planes using a 2D bed-spring model. A 34x34 ball 1mm pitch BGA attaches at the center of a 4" x 4" 3mil thick plane represented by the bed-spring array. The capacitors are located 0.16" from the outside device balls. We modeled 1.5^N array as hypothetical fractional capacitors evenly distributed around the capacitor periphery.

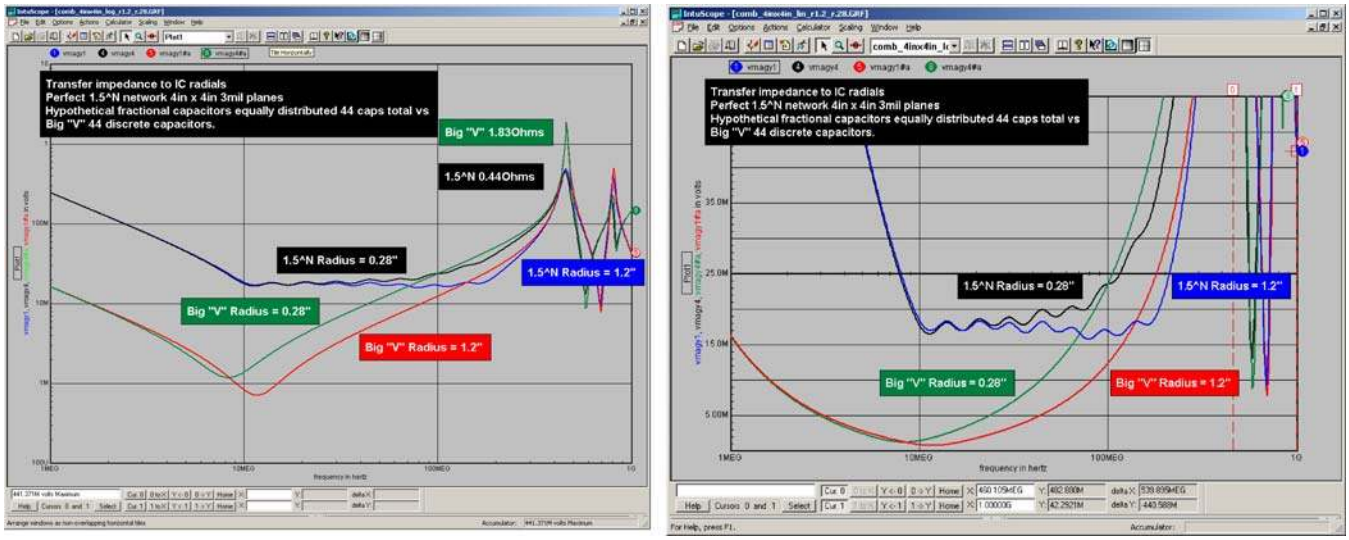

Figure 11, Example 1.5^N Versus Big "V" at IC Attachment

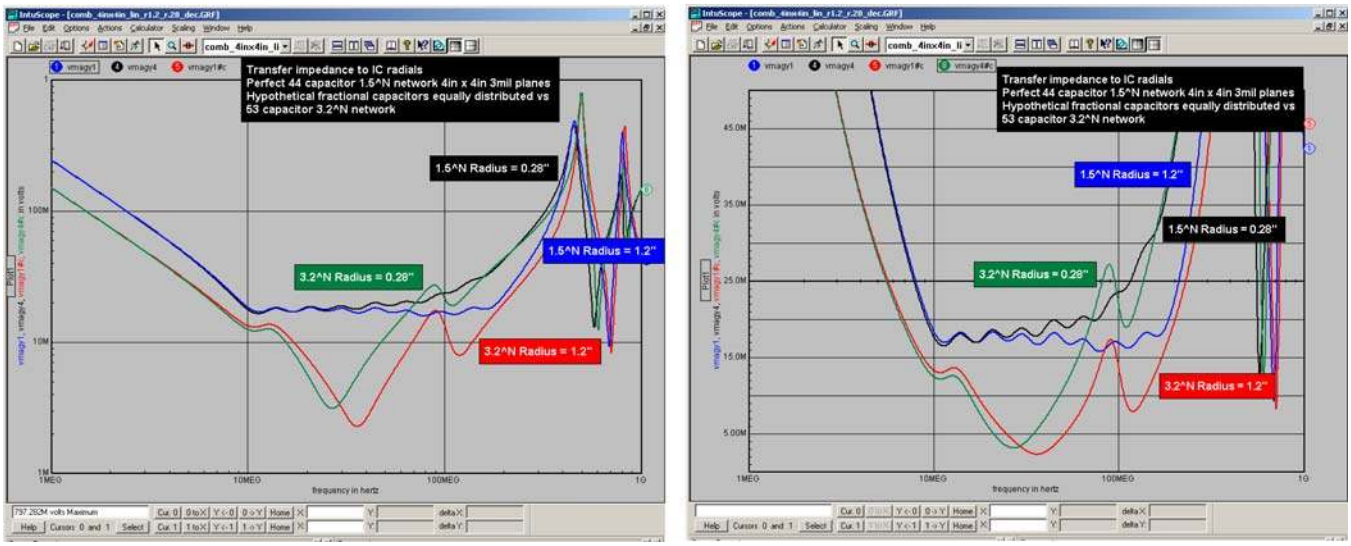

Figure 12, Example 1.5N Versus 3.2N at IC Attachment

The fine multi-pole approach retains its component efficiency advantage best when used with packages that do not crowd the power connections into a center slug configuration and include very thin planes for internal power distribution to the die. Common BGA packages with center slug configurations, such as the 0.28" radius examples shown in Figure 11, diminish the component advantage of fine multi-pole against big "V" when applied with nominal plane separations. Spreading inductance takes an even bigger toll on $X^{3.2}$ networks, all but eliminating any advantage over fine multi-pole, even when using 20% more parts.

The 2D models show that fine multi-pole retains significant advantage in peak impedance, 450mohms in this example, at network/plane PRF versus either: big "V", 1.8 ohms, or $X^{3.2}$, 800mohms.

Transient Responses

Applying a steady stimulus of a 150MHz trapezoidal waveform with 2ns Tr/Tf to the networks above, we should expect that the peak amplitudes are similar. 1D simulation bears out that they are:

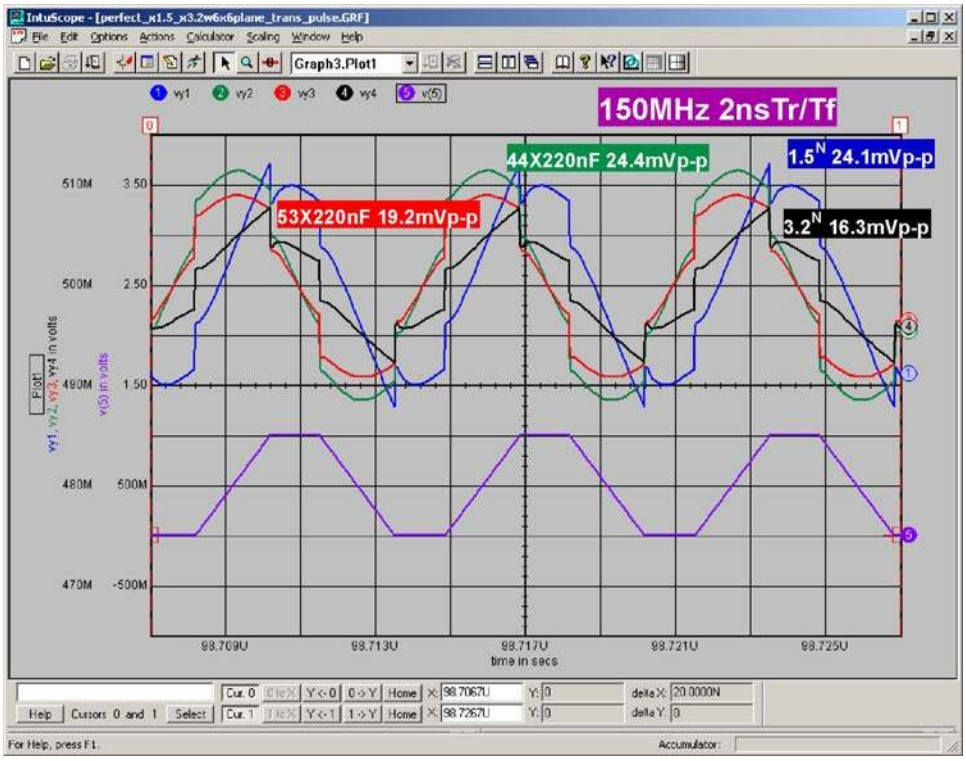

Figure 13, 150MHz Trapezoid, Steady State Transient Response

Figure 13, demonstrates that for this frequency stimulus, the $X^{3.2}$ network affords the lowest noise levels owing to its lower impedance in this region. The remaining networks remain very close in response as predicted by the impedance magnitude plots.

Similarly, if we apply an FM sweep, there are no surprises to the response:

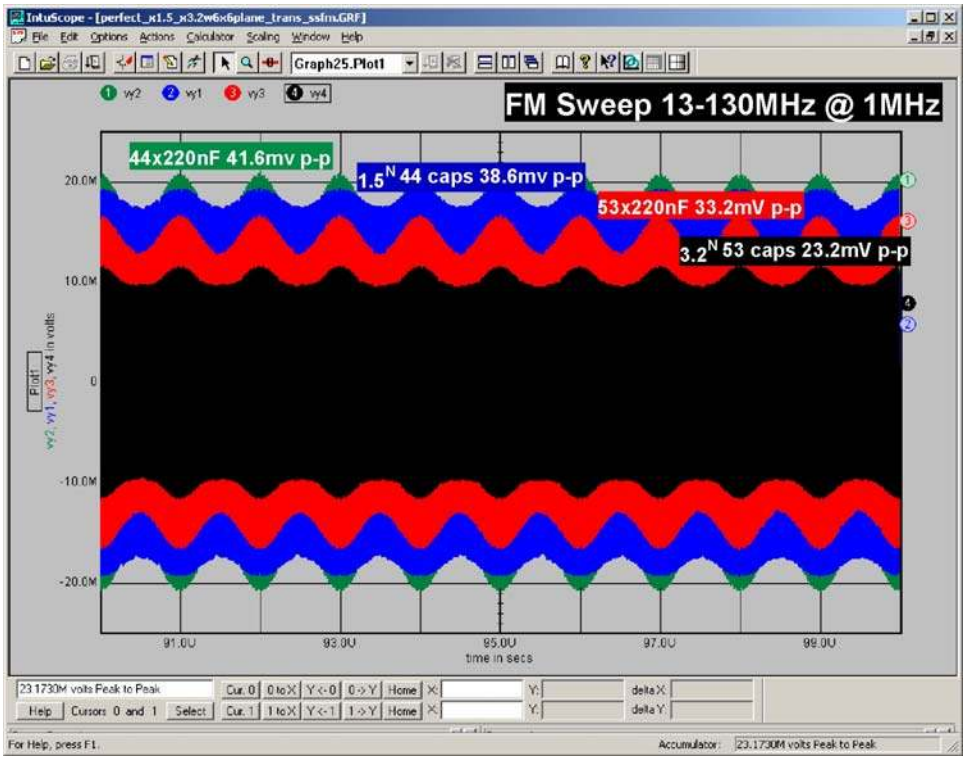

Figure 14, Transient Response, FM Sweep

As expected the low ripple in the many pole network results in the flattest reaction to the FM sweep. Owing to its improved response near the carrier frequency, the 3.2^N network shows the lowest peak amplitude.

As a final measure of transient response, we observe the behavior in response to an isolated pulse.

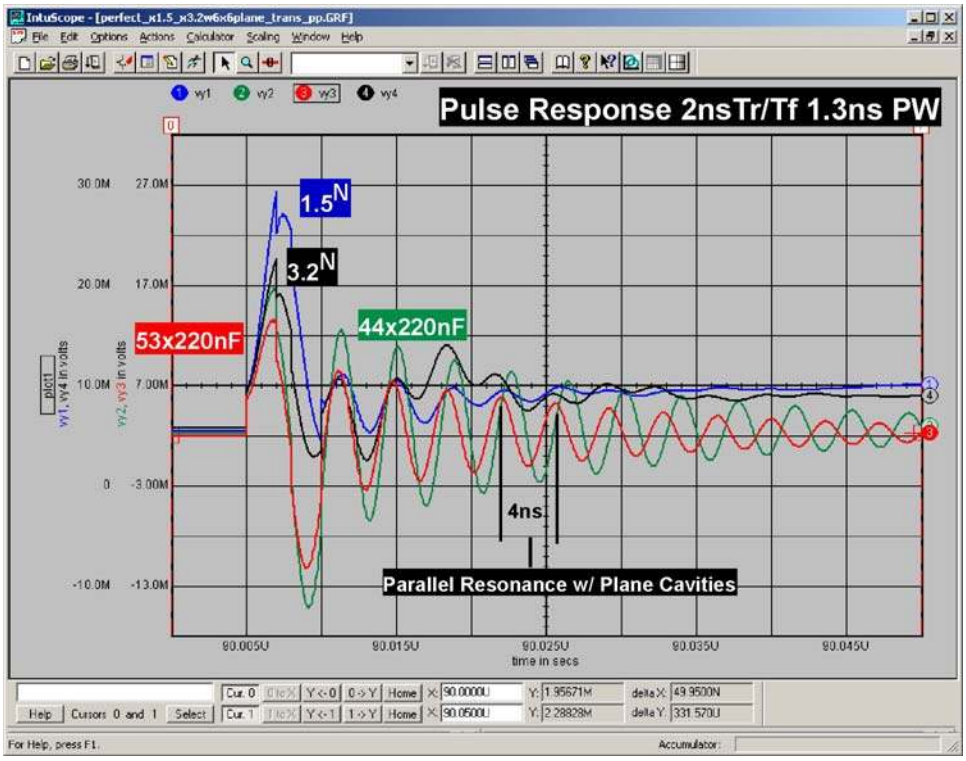

Figure 15, Transient Response, Isolated Pulse

It is apparent from these 1D results that the 1.5^N network has the highest initial amplitude. However, the big "V" implementations exhibit much greater undershoot. The lower damping of the big "V" reflects the much higher peak that occurs at parallel resonance with the plane cavity. Ostensibly, this is out of the power delivery range, but could be a major issue if the bypass network is used for signal return currents. A repetitive signal excitation source such as a "1-0-1-0" DDR2 533Mbps pattern could excite any of these networks into steady oscillation. The situation is far worse for the big "V" networks than the multi / many pole networks.

VRM / Bulk Capacitor Transition

Whereas plane characteristics are fairly easy to define, the VRM / bulk capacitor combination can take on several forms. We concentrate here on the most difficult case, which is a commercial converter using MLCCs as the output capacitors. The resulting device has high Q, which challenges containment of parallel resonance at the transition to the PWB bypass capacitors. Equation 9 dictates that we cross we need substantial phase margin to limit parallel resonant peaking. To achieve this, we need to cross with the converter MLCCs close to their own mounted SRF. This is a parameter that we will most likely have to find by measurement. Once the parasitics have been determined, we can solve for the left hand side of our filter response.

Figure 16 illustrates an example case where despite much higher ESRs, the X^N networks peak badly, while the big "V" networks are well behaved. This should drive home the point that in order to manage peaking, we really are managing the phase margin at the impedance magnitude crossover. Because the X^N networks in this example cross at a much higher frequency where the VRM / bulk phase has reached nearly 90 degrees, the higher ESR of the X^N networks is of little help containing the substantial resonance. The big "V" works well in this example, because despite being at almost –90 degrees phase the impedance magnitude crosses the VRM / bulk capacitors where they are at a low phase angle.

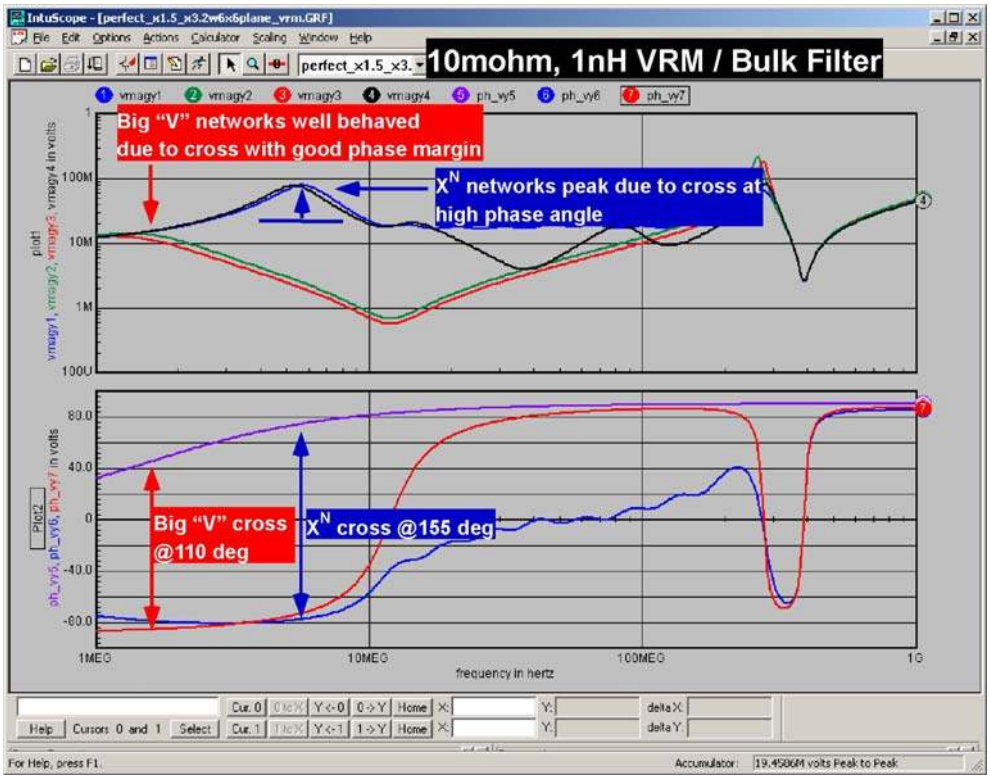

Figure 16, Example w/VRM / Bulk Filter Transition

Given either an X^N or a big "V" network design, and a model for the VRM / bulk network, it is a straightforward matter to check the phase margin and peaking. If sufficient phase margin is not available, then additional, intermediate network(s) will be required to realize an acceptable response.

With the big "V", we have the option of increasing the individual capacitor values until we hit a price point inflection. Presently, this tends to occur at the largest one or two values of capacitance for a given case size, chemistry and voltage rating combination. Since application of the big "V" is built on the premise that parallel resonance with the planes is not a design issue, the sage advice of Dr. Howard Johnson to employ the largest (we suggest the largest value before price inflection) capacitance in a given capacitor case size / chemistry combination is well taken. This approach affords the highest probability that the transition from the VRM / bulk capacitors occurs with sufficient phase margin to meet the impedance profile at little or no cost impact.

For X^N networks, limitations in available capacitor values dictate practical combinations of new filter networks.

It is observable that for large capacitance MLCCs, the component cost is fairly proportionate to capacitance. Some modest gains in placement and via drilling costs are to be had by using the largest capacitors, ie fewest additional poles necessary to satisfy the impedance profile.

ESR Sensitivity

X^N networks exhibit increased transfer impedance for ***lower*** component ESR values. The sensitivity rises with larger values of X.

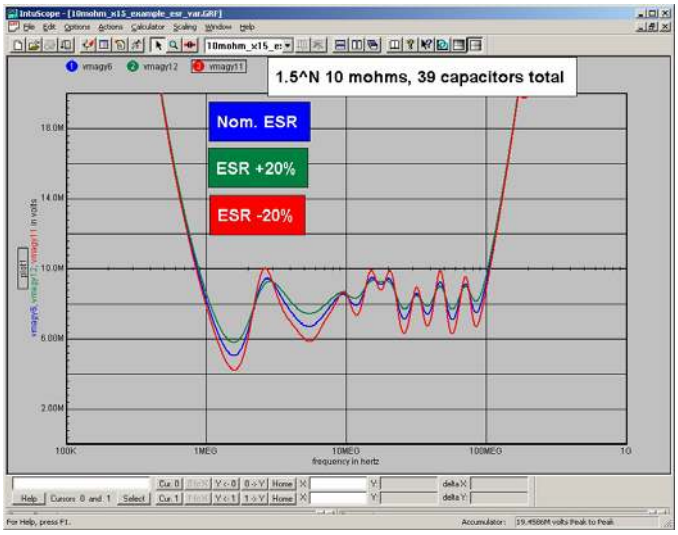

Figure 17, Example ESR Sensitivity 1.5^N Network

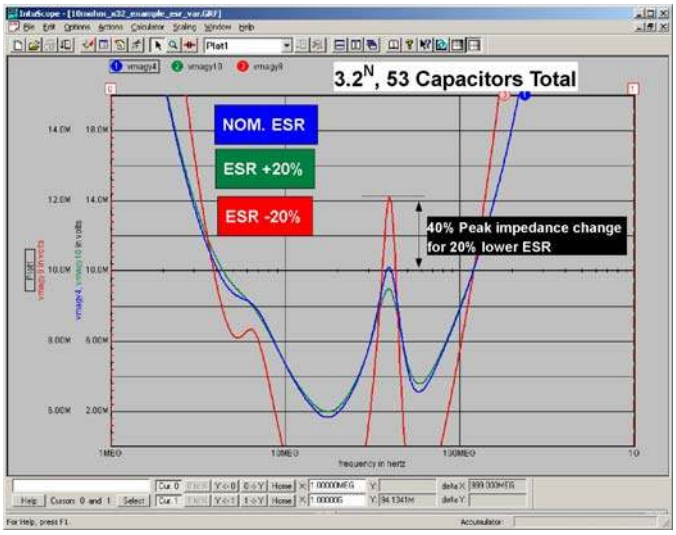

Figure 18, Example 3.2^N Network

On the other hand, big "V" networks demonstrate little sensitivity to ESR:

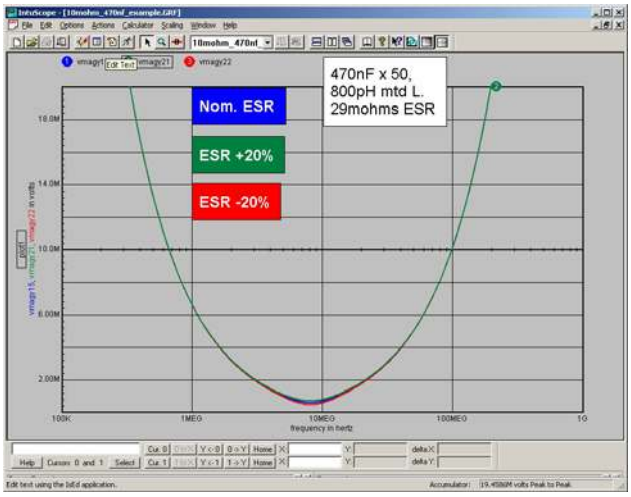

Figure 19, Big "V" ESR Sensitivity

Mis-Stuffed Component Sensitivity

X^N networks depend upon carefully matched quantities of each capacitor value. It should come as little surprise that minor stuffing errors cause big impedance perturbations.

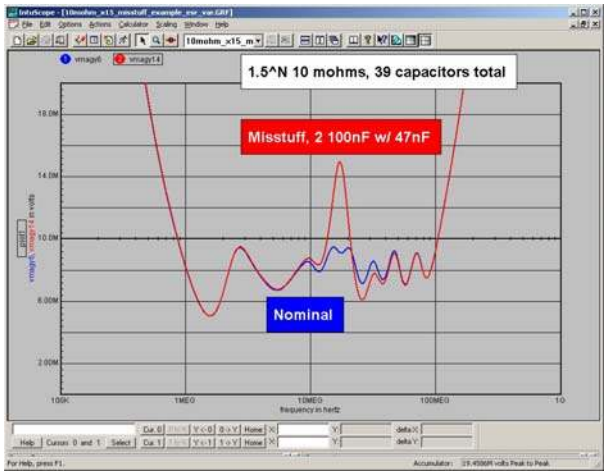

Figure 20, 1.5^N 5% Mis-stuffs

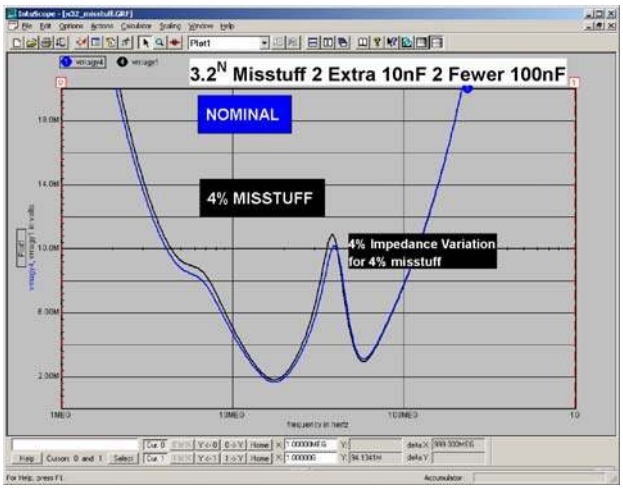

Figure 21, 3.2N 4% Mis-stuffs

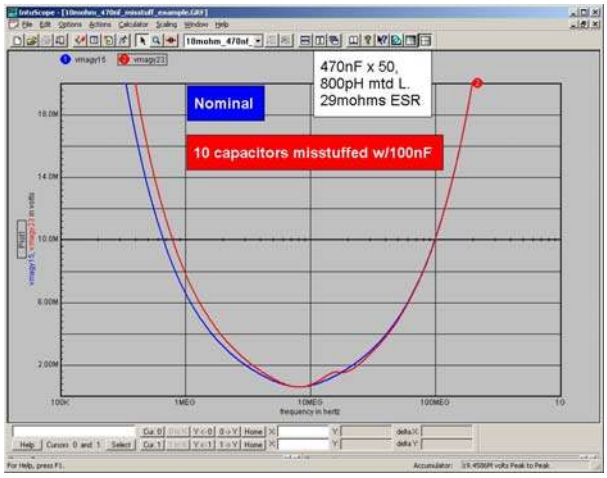

Figure 22, Big "V" w/20% Misstuffs

The good news for each multi-pole case is that each capacitor has a clear frequency signature. The bad news is that board QA really needs a frequency sweep. If an error is found, most likely all capacitor locations for the offending value need to be removed and replaced.

3.2^N networks have much greater redundancy than 1.5^N networks and demonstrate a greater tolerance to misstuffed values. While big "V" networks enjoy even greater redundancy, and even 20% misstuffs have little effect on performance.

Open or Missing Component Sensitivity

Sensitivity to open components is worst in the X^N networks at the low frequency end where the number of components in a given value is few. Smaller X suffers more than larger X. Big "V" networks show very little sensitivity to open or missing components.

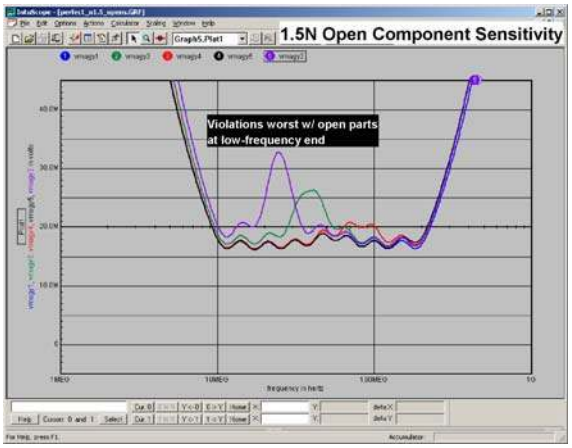

Figure 23, 1.5^N Open Component Sensitivity

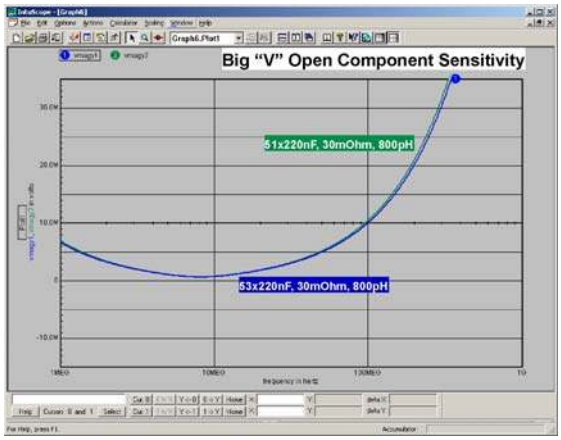

Figure 25, Big "V" Open Component Sensitivity

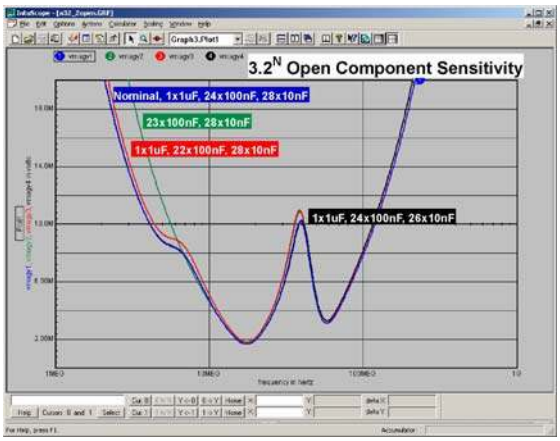

Figure 24, 3.2^N Open Component Sensitivity

Network Synthesis

Big "V" network synthesis is straightforward:

1. Determine the high frequency cut-off frequency and impedance.
2. Translate this to an equivalent inductance.
3. Find device count by dividing the target inductance into the mounted inductance of a single capacitor.
4. Map the parallel capacitance response against the VRM / bulk capacitor response, and determine if adequate phase margin is available to prevent excessive peaking.
5. If peaking is excessive, insert an additional intermediate capacitor value with SRF approximately at the impedance crossover point.
6. Scale quantity until the impedance profile is satisfied.

X^N Synthesis is not much more difficult:

1. Select a transition frequency from the VRM / bulk capacitors that affords at least 45 degrees phase margin.
2. Determine the high frequency cut-off and target impedance.
3. Estimate the number of capacitor values N_{MAX} as:

Equation 16 $N_{VAL} = ROUND\_UP(\ln(F_{HF\_CUT\_OFF}/F_{LF\_CROSS})/ \ln(X)$

4. For each N, from 0 to N_{VAL-1} determine an initial capacitor count based on K2 = -0.5:

Equation 17 $COUNT\_INITIAL_0 \approx K4 * ESR_0 / |Z_{TARGET}|$

Equation 18 $COUNT\_INITIAL_N \approx K4 * X * ESR_N/|Z_{TARGET}|$

Note that K4 decreases with decreasing X, and increases with increasing mounted capacitor inductance.

5. For each N, from 0 to $N_{VAL} - 2$ correct the capacitor count to reflect K2 > -0.5:

Equation 19 $K_{COUNT\_COMP\_EXP(N)} \approx K5 - (\ln(ESR_{CAP(N+1)}/ ESR_{CAP(N)})/ \ln(X))$

Equation 20 $COUNT_N \approx ROUND\_UP(COUNT\_INITIAL_N^{KCOUNT\_COMP\_EXP(N)})$

6. Verify the synthesized network performance in SPICE. Adjust as necessary.

Note that for small values of X, sensitivity to count quantization can be quite high.

Summary

Bypass filter design is as much about insuring adequate phase margin at each frequency transition as any other task. The many pole approach developed by Sun researchers maintains phase margin not only at the network extremes, but across the entire filter band as well. This provides the greatest advantage of lowest amplitude impedance at parallel resonance with the PCB plane cavity. This may be a significant EMC issue, as well as where signal return images reference multiple voltage planes.

Device ESR greatly affects phase looking up in frequency and so affects transitions between adjacent MLCC networks, and the transition from the discrete capacitor networks to the PCB plane cavities. However, MLCC ESR is of little concern down in frequency at the transition from the VRM / bulk capacitors. At the bulk transition, we need sufficient capacitance from the higher frequency network so that the transition occurs at a low enough frequency that the VRM / bulk capacitor network phase is well below 90 degrees. This is usually a significant advantage for big “V” networks.

Total component count, and component cost advantage swings between the multi / many pole approaches and the single capacitor value approach depending on the combined circumstances of the PWB target impedance and VRM / bulk capacitor configuration.

Viable bypass capacitor networks may be synthesized by any of the three popular techniques using straightforward spreadsheet equations backed by SPICE simulations. 1D network design fails to account for important spatial effects. 2D simulations should be used to verify and adjust any design. No matter which technique is used, low mounted inductance reduces required component account across the frequency spectrum, not just at mid to high frequencies.

Of the three basic approaches, the most appropriate method depends on how one weights various design and manufacturing criteria. Table 2 summarizes a number of these criteria and serves to illustrate that each method has specific strengths and limitations relative to the other two:

Criteria	1.5^N	3.2^N	"V"	Comments
Plane cavity parallel resonance with discrete capacitors	▲	▲-	▼	Matters when the application uses bypass capacitors for signal returns
Total parts count, optimized	?	?	?	Depends on VRM / Bulk cap configuration, target impedance, bandwidth and spatial power delivery requirements
Total parts cost, optimized	?	?	?	
Transient pulse, initial peak	▼	▲	—	Rings at discrete capacitor to plane resonance
Transient single pulse, p-p	—	▲	▼	
Transient pulse decay time	▲	—	▼	
Response close to HF cut-off	▲-	▲	▼	3.2^N capable of deep SRF closest to cut-off
Low and mid frequency response	▼	—	▲	Is the noise profile flat?
Manufacturing complexity	▼	▼+	▲	Typical twelve values versus three versus one X^N boards S/B swept with SA or VNA after assy' to insure proper stuffing
Immunity to ESR variation	▲-	▼	▲	Continuous supply monitoring critical for 1.5^N
Immunity to stuff errors	▼	▲-	▲	1.5^N requires removal and replacement of many components in the event of one or a few mis-stuffs. Errors in larger capacitor values cause more severe impedance impacts.
Immunity to open components	▼	▲-	▲	

Table 2, Summary, Bypass Method Trade-offs

.

Conclusions

Each of the three common bypass capacitor methods: many-pole, capacitors on decade values, or single bypass capacitor value enjoy strengths and suffer weaknesses compared to the alternative methods. Properly applied, each method can yield well-performing networks. A dispassionate review of design and manufacturing criteria will aid the astute designer in the selection of a preferred method, and the wisdom to apply either of the alternatives under the appropriate circumstances.

Fine multi-pole can reduce required parts counts where either the PCB, and/or the IC package use very thin plane cavities. Fine multi-pole also yields the lowest impedance peaks at the capacitor to board PRF. Fine multi-pole has better resilience to ESR variation than capacitors on decades spacing, but correct population particularly of the larger capacitor values is critical.

Capacitors on decade values can provide some impedance advantages over fine multi-pole at the cost of somewhat higher peak impedance at PRF. They do not improve component count over fine mult-pole at the high frequency end, fail to provide the bulk capacitor transition advantages of big "V" at the low-end. They also suffer the highest sensitivity to ESR variations. For equal capacitor counts they can selectively stretch bandwidth compared to big "V".

Big "V" offers great simplicity and manufacturing tolerance. In many cases, big "V" delivers high frequency IC power equally as either X^N method using the same number of components. Do to its wider bandwidth, big "V" will often result in fewer total components, and lower cost when the transition to the bulk capacitor network is taken into account. Big "V" however exhibits a much higher board PRF than an equivalent performance X^N network.

References

1. Larry D. Smith, et-al " Power Distribution System Design Methodology and Capacitor Selection for Modern CMOS Technology" IEEE Transactions on Advanced Packaging, Vol. 22, No. 3, August 1999, pp284-291
2. Howard Johnson, PH.D. Martin Graham, PH.D. "High-Speed Digital Design A Handbook of Black Magic" Prentice Hall, 1993, pp. 281-293
3. William J. Dally, John W. Poulton, "Digital Systems Engineering" Cambridge University Press, 1998, pp. 247-256.
4. Brian Young, "Digital Signal Integrity", Prentice Hall, 2001, pp 412-420.
5. Tanmoy Roy, Larry Smith, John Prymak "ESR and ESL of Ceramic Capacitor Applied to Decoupling Applications", EPEP Proceedings 1998.
6. Larry D. Smith, Jeffery Lee, "Power Distribution System for JEDEC DDR2 Memory DIMM" EPEP Proceedings, 2003.
7. Steve Weir, Scott McMorrow, "High Performance FPGA Bypass Networks", DesignCon Proceedings 2005.

[1] In the ideal case where K2 from Equation 1 equals -0.5, this relationship would be exact. As K2 is in reality higher, Q creeps upward with each stage, and capacitor count grows more quickly than X^N.