

DesignCon 2006

Frequency Domain Target Impedance Method for Bypass Capacitor Selection for Power Distribution Systems

Larry D Smith, Altera Corporation

lsmith@altera.com

408-544-7822

Abstract

There has been much discussion in the industry on how to choose bypass capacitors for power distribution systems. This paper discusses the Frequency Domain Target Impedance Method. A target impedance is defined. The power system should meet target impedance across a broad frequency range, from DC up to the highest frequency of interest. Simulation is used to select capacitor values and quantities from a menu. The number of capacitors required to meet the target impedance up to a corner frequency is estimated by a simple formula. High impedance resonances are avoided by providing sufficient system damping. The relationship of PCB capacitors to the VRM at low frequency and the mounted die at high frequency is discussed. Capacitor *ESR* is an important consideration. The Frequency Domain Target Impedance Method is compared to the Big V and Decade Methods for choosing bypass capacitors. It is found to have superior performance and cost characteristics.

Author Biography

Larry D Smith received the BSEE degree from Rose Hulman Institute of Technology in 1975 and the MS degree in Material Science from the University of Vermont in 1983. After joining IBM in 1978, he worked in the areas of reliability, characterization, failure analysis, power supply and analog circuit design, packaging and signal integrity. He worked at Sun Microsystems from 1996 to 2005 where he did much development work in the field of power integrity. In August of 2005, he joined Altera Corporation where he is concentrating on SSN noise, power and signal integrity. He has 13 patents and has authored numerous Journal and conference papers.

Introduction

Bypass capacitors, otherwise known as decoupling capacitors, are an important part of the Power Distribution System (PDS). There has been much controversy over how to choose the optimum set of capacitors to accomplish a reliable PDS. The primary figures of merit for a PDS include the noise performance, cost, reliability and surface area consumed by the capacitors.

The primary components of the PDS include the voltage regulator module (VRM), bulk capacitors, high frequency ceramic capacitors, PCB power planes, capacitor mounting pads and vias, mount for the electronic package (load, power consumer), package capacitors, package power planes, die mount and internal die capacitance. Many of these components are shown in Figure 1, roughly in the frequency range where they are effective.

This paper concentrates on the Frequency Domain Target Impedance Method (FDTIM) for selecting bypass capacitors that will work with the rest of the system to provide clean power to the load, usually one or more silicon die. The method has been previously discussed in [1]. A key concept is the determination of a target impedance. By meeting the target impedance from DC up to the highest frequency of interest at each level of assembly, the optimum compromise between performance, cost, reliability and board area is reached.

Target Impedance

As part of the die design process, engineers assume a nominal power supply voltage with a tolerance, often plus or minus 5%. Circuits are simulated to guarantee timing and performance specifications assuming that the power supply at the circuit terminals meets a well defined specification. This simulation usually involves SPICE analysis with ideal voltage sources at three ‘corners’ to provide power at nominal, nominal – 5% and nominal +5% voltage conditions.

In the real world, there are no ideal SPICE power supplies because real voltage sources have series source impedance. This series impedance gives rise to ‘load regulation’ as the magnitude of the current drawn by the load changes during normal operation. The cost of a power supply is inversely proportional to the series impedance, with an ideal zero impedance power supply being infinitely expensive! The optimum power supply is one that delivers the specified power requirements at the minimum cost. By knowing the nominal power supply voltage, maximum and minimum load currents and the allowable voltage tolerance, the target impedance is easily calculated from Ohms law.

$$Z_{\text{target}} = \frac{V_{dd} \times \text{tolerance}}{I_{\text{max}} - I_{\text{min}}} = \frac{1.0 \text{ V} \times 0.05}{100 \text{ A} - 50 \text{ A}} = 0.001 \text{ Ohms} = 1 \text{ mOhm} \quad (1)$$

In this example, the nominal power supply is 1.0V; the tolerance is 5%; and the maximum and minimum load currents are 100A and 50A respectively. A 1V power supply with 1 mOhm of source impedance will vary by 0.05V when the maximum and

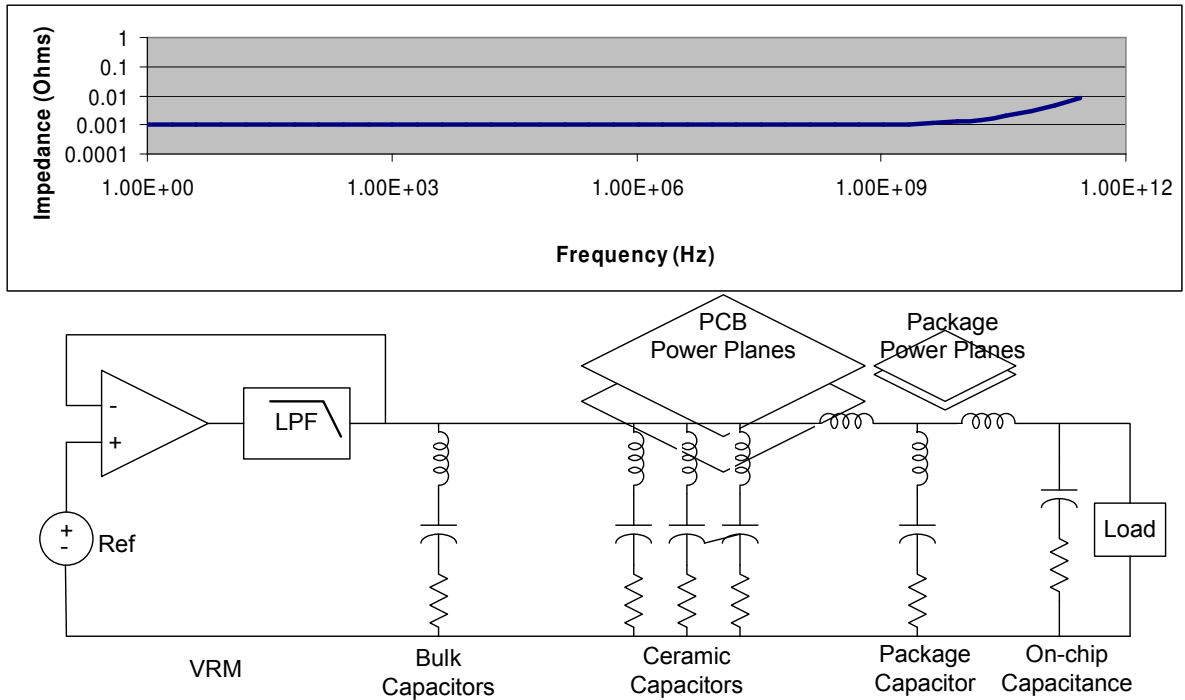


Figure 1: Components for Power Distribution System (PDS) design. The impedance of the PDS should be flat across a broad frequency range.

minimum currents are drawn by the load. The difference between maximum and minimum load current is often called the transient current. Both the magnitude and rise (fall) time of the transient is important. The magnitude determines the target impedance; the rise time determines the upper corner frequency where the target impedance should be met.

It may appear that this calculation is off by a factor of 2 because plus and minus 5% voltage tolerance is allowed. But the regulation loop of a power supply VRM will bring a 5% low voltage, sensed at the load, back to nominal voltage after the VRM time constant has past. When the load goes away, the voltage at the load will bounce to plus 5% until another VRM time constant has past, thus using up the entire +/- 5% tolerance.

The target impedance calculation applies at all levels of assembly including a single circuit, block of circuits on a die, full die, package, power plane on a printed circuit board (PCB) that supplies several loads, the DC source for VRMs (often 12V or 48VDC) and even at the AC system supply level.

The target impedance is useful for sizing the amount of capacitance necessary at each level of assembly to store sufficient charge and energy for the load. When a fast load draws a power transient, charge stored in small capacitors near the load supply the current for a short period of time but are soon depleted. Current must then come from slower time constant capacitors at the next level of assembly. Typical PDS charge storage areas include on-chip capacitors, on-package capacitors, PCB ceramic and bulk

capacitors and capacitors associated with the input and output of the VRM. Each level of assembly provides transient current long enough for the time constant of the next level to expire so that energy can begin to come (or stop coming) from the next stage.

A PDS that meets target impedance across the entire frequency span has sufficient stored charge to supply clean power at all frequencies. On the other hand, a power system that is higher than the target impedance in any frequency band will usually fail to meet the specified voltage tolerance under some load condition. If the measured impedance of a PDS is substantially below target impedance in any frequency range, the cost of the PDS can usually be reduced while still meeting the power specifications. It is referred to as a 'target' because it gives the best cost/performance solution for the PDS.

Impedance in the frequency domain

The target impedance should be met not only at DC but at all frequencies up to some corner frequency, f_c . There are mechanisms at the die and system level that can cause power supply transients in all frequency bands, not just the clock frequency. The circuits on the die have rise times on the order of 10 pSec. Using the common formula that relates frequency content to rise time $f_c = 0.35/t_{rise}$, the largest frequency of interest for an individual circuit is 35 GHz. A die operating with a clock frequency of 1 GHz can suddenly begin to draw (or cease drawing) current from the power supply in about 1nSec giving a corner frequency of about 350 MHz. Integrated circuits can suddenly draw transient current after a key stroke or cease to draw current when operations stall while waiting for data from memory. It may take 40nSec to retrieve data from DRam and a die could repetitively draw pulses of power current in 100nSec cycles putting a 10 MHz load on the PDS. Or, the CPU could be waiting for data from a hard drive with a 1 mSec time constant leading to kHz loading on the PDS. It is possible for customer code to cause repeating current transients at virtually any frequency from DC to 35 GHz therefore the PDS should meet target impedance in all frequency bands in the respective levels of assembly. The circuits on the die should look out and see an impedance that is close to target impedance throughout the entire frequency range.

Simulation is used to select the best mix of capacitors. The impedance can be measured by using a 1 amp current source in the die position of the circuit. The voltage across the current source is numerically the same as the self impedance of the PDS measured from the die position. Voltages at other positions of the circuit are numerically the same as the trans-impedance which is the voltage across any two nodes of the circuit divided by the current forced at some other position of the circuit. By selecting optimum values and quantities of bypass capacitors to be placed at various locations of the circuit, the PDS can be designed to have a relatively flat impedance profile across a broad frequency range. Figure 1 shows a 1 mOhm target impedance up to the corner frequency associated with the rise time of a silicon gate along with the components that are affective in each frequency range.

The major difficulty with PDS design is that the system is built from a network of inductances and capacitances that attempt to mimic the flat impedance of a 1 mOhm resistance. A conductor always has parasitic inductance which has a plus 20dB per

Cap Value	Size	Dielectric	Measured Value	Units	ESR (mOhms)	L interanl (nH)	L mount (nH)	SRF (Mhz)	Q	ESR divided by 1 mOhm	Meet 1 mOhm Ztarget
100uF	1812	X5R	80.3	uF	1.8	2.112	0.600	0.341	0.7	2	2
47uF	1210	X5R	42.1	uF	1.9	1.487	0.600	0.537	1.1	2	3
22uF	1210	X5R	17.7	uF	2.5	1.300	0.600	0.867	1.3	3	7
10uF	0805	X5R	7.26	uF	3.6	0.773	0.600	1.60	1.6	4	9
4.7uF	0805	X5R	4.12	uF	4.2	0.544	0.600	2.32	2.1	4	5
2.2uF	0805	X5R	1.98	uF	6.1	0.413	0.600	3.55	2.2	6	8
1.0uF	0603	X5R	0.79	uF	9.1	0.391	0.600	5.69	2.3	9	12
470nF	0603	X5R	404	nF	13	0.419	0.600	7.85	2.3	13	16
220nF	0603	X7R	172	nF	19	0.438	0.600	11.9	2.3	19	28
100nF	0603	X7R	75	nF	29	0.443	0.600	18.0	2.3	29	30
47nF	0603	X7R	39	nF	38	0.451	0.600	24.7	2.4	38	40
22nF	0603	X7R	17	nF	64	0.492	0.600	36.6	2.1	64	53
10nF	0603	X7R	8.9	nF	80	0.518	0.600	50.4	2.4	80	60
Totals										273	273

Table 1: Discrete MLC capacitors that are useful for PCB bypass. Measured values for capacitance and inductance are shown together with the series resonant frequency and Q for 600 pH mounting inductance

decade slope on a log-log impedance plot. Intentional capacitance from discrete capacitors and parasitic capacitance from power planes has a slope of minus 20dB per decade. Resonance occurs where the two slopes cross. Parallel combinations of capacitance and inductance form impedance peaks while series combinations form impedance dips. For power loss reasons, resistance is often designed out of power systems and can lead to high Q resonances if the crossings are not well controlled. A major objective in choosing decoupling capacitors is to insure that resonant peaks that exceed the target impedance are avoided. This is done through the use of accurate frequency domain models for each of the PDS components.

PCB Bypass Capacitor Selection

Much board space in modern computer systems is dedicated to the power distribution components, mostly bypass capacitors. There is cost associated with the components, assembly process, vias and the PCB area occupied by the capacitors. It is important to get the most benefit out of each mounted capacitor in order to minimize cost. Important considerations include capacitance value, equivalent series resistance (*ESR*), internal inductance and mounting inductance.

Discrete PCB capacitors are used to cover the frequency range where the VRM leaves off up to where the inductance of the mounted electronic package dominates. This is typically the 10kHz to 100MHz band, about 4 decades. The upper part of the frequency range is covered by multi layer ceramic (MLC) capacitors and the lower part is covered by aluminum electrolytic, tantalum oxide or tantalum polymer bulk capacitors. Advancements in VRM technology have increased the bandwidth and response times to the point where the bulk capacitors may not be necessary but this usually involves a more expensive VRM. Table 1 gives the size, dielectric type and *ESR* associated with several ceramic capacitor types that are useful for bypassing.

The Frequency Domain Target Impedance Method for PDS design involves choosing the number of capacitors of each value such that the parallel combination of capacitors approximates a flat line at the target impedance. The approximate number of capacitors




Figure 2: Simulated impedance for several bypass capacitor values, a) single capacitors, b) quantity to reach 1 mOhm. Capacitors are selected from this menu to build a power distribution system. The vertical scale of Volts should be interpreted as Ohms because of the 1 amp current source used in simulation.




Figure 3: Model to hardware correlation for several ceramic capacitors.

for each value can be calculated by dividing the target impedance into the *ESR* of that capacitor.

Figure 2a shows the simulated impedance of the capacitors of Table 1 using the transmission line model topology for the capacitors discussed in [2]. Figure 2b shows simulation results for the number of capacitors in parallel needed to reach 1 mOhm as calculated in Table 1. Three capacitor values are used per decade: 10uF, 22uF, 47uF, etc. Each capacitor curve has three distinct portions: capacitance at -20dB/decade, inductance at +20dB/decade and the bottom of the curve at the *ESR* value for the capacitor. This simulation includes only the internal inductance for capacitors (does not include mounting inductance). The family of curves shows that decreasing valued capacitors have a series resonance at increasing frequency. The *ESR* of smaller valued capacitors is higher than that of lower valued capacitors. This is the menu of capacitor components that may be used in a PDS to meet a target impedance within this frequency range.

Simulation parameters for capacitor models are chosen for best fit with capacitors measured using two port techniques with a VNA [3,4]. Figure 3 shows model to hardware correlation.




Figure 4: a) Simulation results for parallel capacitors calculated to reach 1 mOhm. b) Capacitors chosen to stay below 1 mOhm target impedance. Quantities of each capacitor value are given in Table 1.

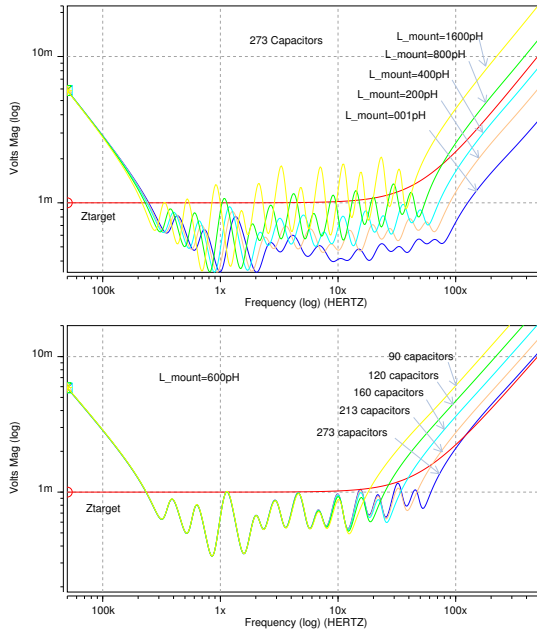


Figure 5: a) Simulation results for capacitors with several values of mounting inductance. b) A reduced number of capacitors is able to meet target impedance but only up to a lower corner frequency.

Figure 4 shows simulation results for 273 total capacitors in parallel compared to a 1 mOhm target impedance. The specific quantities of capacitors are given in the column of Table 1 “Quantity for 1 mOhm.” The simulated impedance drops below 1 mOhm because of the influence of capacitors at near by resonant frequencies. Figure 4b shows the simulation results for the same number of capacitors but value quantities are chosen to meet the 1 mOhm target impedance in the lower frequency range and are given in the final column of Table 1.

Capacitor Sizing from Target Impedance and Corner Frequency

Fifty MHz has been chosen as the corner frequency, above which the target impedance will not be met at the PCB level of assembly. It is chosen from knowledge of the mounted inductance of the load. There is little point in pushing the low impedance on the PCB much more than 2x the frequency supported by the package mounting inductance. The corner frequency is the frequency where the target impedance equals the parallel combination of all capacitors which have become inductive.

$$Z_{\text{target}} = \left| j\omega_c \frac{L_{\text{cap}} + L_{\text{mount}}}{n} \right| = 2\pi f_c \frac{L_{\text{cap}} + L_{\text{mount}}}{n} \quad \text{and} \quad f_c = \frac{n \cdot Z_{\text{target}}}{2\pi(L_{\text{cap}} + L_{\text{mount}})} \quad (2)$$

where n is the number of capacitors in parallel, L_{cap} is the internal capacitor inductance and L_{mount} is the mounting inductance. After all the capacitors have become inductive, the location of the rising slope is simply the impedance of all the inductances in parallel.

Figure 5a shows curves with several values of mounting inductance added in series with each capacitor: 200, 400, 800, 1600 pH. This is the range expected depending on whether great priority or little priority is given to achieving low mounting inductance in the PCB geometries. As shown in the figure, the resonant minimum of each individual capacitor moves up and to the left with increasing mounting inductance. The bottoms of the individual capacitor curves are initially rounded, then become sharper as the mounting inductance is increased. As expected, the corner frequency of the entire set of capacitors in parallel moves to the left with increased mounting inductance.

Figure 5b shows simulation of a reduced number of capacitors by eliminating the highest frequency capacitors first. The same target impedance is met but only up to a reduced corner frequency.

From the examples and figures above and a number of completed designs, it is estimated that the number of ceramic capacitors required to meet a target impedance in the mOhm range with a corner frequency in the 50 MHz range is

$$n = k \frac{f_c \cdot (L_{cap} + L_{mount})}{Z_{target}} \quad (3)$$

where k is about 6, average L_{cap} is $0.4nH$, L_{mount} is the mounting inductance in nH , Z_{target} is the target impedance in $mOhms$ and f_c is the corner frequency in MHz . This equation is for capacitors with *ESRs* that are commonly available today.

ESR Considerations

It is apparent that there is a relationship between capacitor *ESR* and the number of capacitors required to meet a target impedance. The initial estimate for an individual capacitor quantity was ESR/Z_{target} and it would appear that a lower *ESR* would result in fewer capacitors required. But the low *ESR* that creates deep resonant dips may also create high resonant peaks unless the Q (mounting inductance) is well controlled.

The *ESR* in commonly available X7R and X5R ceramic capacitors combined with reasonable (600 pH) mounting inductance gives a $Q = \omega L / R$ of about 2 as shown in Table 1. Three capacitor values per decade, as simulated above, appears to be optimal. If *ESRs* were lower, more capacitor values per decade would be required to keep the resonant peaks from getting out of hand. If *ESRs* were higher, fewer capacitor values per decade would be required but a greater quantity of each value would be required in parallel to reach the target impedance. The FDTIM will work in any case.

There have been proposals in the industry for a controlled, higher *ESR* for ceramic capacitors [5]. This has merit because it would reduce the number of different part number capacitors on a PCB and reduce manufacturing complexity. But what would the optimum value of *ESR* be? It would most likely turn out that the optimum *ESR* for several different products would be several different *ESRs* for the same capacitance value (i.e. the optimum *ESR* for a 1 μF capacitor would be 10 mOhm, 20 mOhm and 40 mOhm

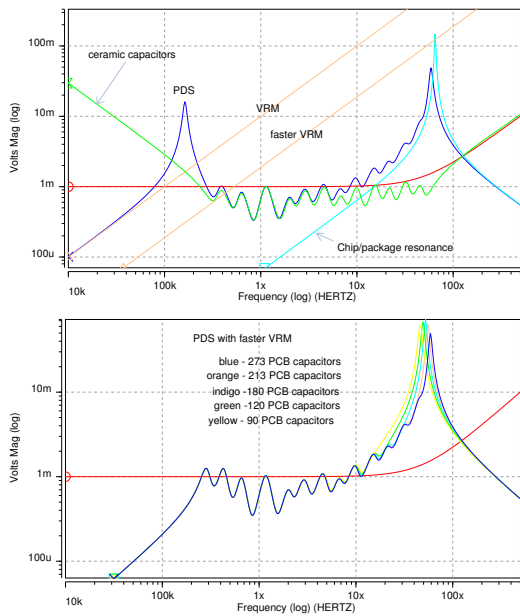


Figure 6: a) PDS impedance peaks due to VRM at low frequency and chip/package resonance at high frequency. b) Faster VRM improves low frequency performance. Chip/package resonance is not improved by any combination of PCB capacitors.

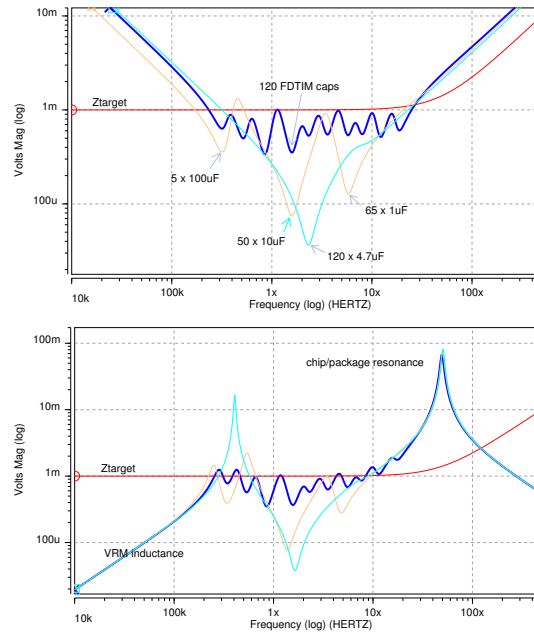


Figure 7: 120 Capacitors selected by Deep V and Decade method compared to FDTIM. a) PCB capacitors by themselves. b) Same capacitors together with VRM and chip.

for 3 different products), leading to multiple part numbers for each capacitance value. This would be difficult to manage from a components and inventory standpoint.

Several years ago the metallurgy for MLC capacitors was Silver-Palladium. *ESR* values varied by a factor of 4 from supplier to supplier and sometimes with in a single lot from a single supplier making it very difficult to simulate with SPICE models and obtain consistent measured results from product boards. The industry has moved to Nickel base metal resulting in very consistent *ESR* from supplier to supplier for each capacitor value, and somewhat *ESR* lower than the previous metallurgy.

It is the opinion of this researcher that the capacitor suppliers should develop the most consistent and cost effective process possible for manufacturing ceramic capacitors. Today's X7R and X5R capacitors work well. Good PCB design techniques enable a mounting inductance that gives a Q below 3, which is manageable from a parallel resonance standpoint with 3 capacitor values per decade. Any attempt to increase *ESR* to a specified value may result in undesirable cost, complexity, size or inductance for the capacitor and complexity in component selection and inventory.

Problems at High and Low Frequency

FDTIM results in an optimal set of ceramic capacitors that meet a target impedance across a broad frequency range. The low end of this frequency range must transition to the VRM output impedance which is usually inductive. The chip/package resonance [6] is at the high end. Figure 6a shows a system level simulation including a VRM and die mounted in an inductive package. A moderately fast VRM has an output impedance that crosses the target impedance at 100kHz . The inductive slope crosses the PCB capacitance slope at about 200kHz creating a major impedance peak. Also shown in the diagram is a faster VRM that crosses the target impedance at about 500kHz . Figure 6b shows that the VRM peak is virtually eliminated by using the faster VRM.

The chip/package resonant peak is more difficult. In this case, the package mounting inductance is 10pH and the die capacitance is 600nF , reasonable values for modern micro processors. The full 273 PCB capacitors cause the impedance of the PCB to be resistive at resonance. This provides some level of damping and reduces the peak by about a factor of 2 in figure 6a. Figure 6b shows the same peak with the higher frequency PCB capacitors removed and little change in the resonant peak. The additional capacitors give only a marginal improvement in the peak. There is little that can be done to reduce chip/package resonance on the PCB. The chip/package resonant peak must be managed inside the package through the use of package capacitors and low inductance mounting structures.

Compare FDTIM to Other Methods

A “Big V” method for choosing bypass capacitors has been proposed [7]. The concept is to use just one value of capacitance to form a “V” shaped impedance in the frequency domain. Another common method is to pick one capacitance value per decade, i.e. $100\mu\text{F}$, $10\mu\text{F}$, $1\mu\text{F}$, 100nF , etc. Both of these methods can lead to impedance peaks where an inductive slope crosses a capacitive slope. Figure 7 a shows PCB capacitors for the Big V and Decade methods compared to FDTIM. The capacitance at low frequency and the inductance at high frequency is approximately the same for the 120 capacitors used for each method. The Big V has $120 \times 4.7\mu\text{F}$. The decade method has $5 \times 100\mu\text{F}$, $50 \times 10\mu\text{F}$ and $65 \times 1\mu\text{F}$. The FDTIM method uses the first 120 capacitors of Table 1. All three methods hold the PCB impedance near the target impedance from about 300kHz to 30MHz .

Figure 7b shows system simulation with the VRM and chip/package resonance. The big V method is particularly vulnerable at the VRM crossing. There is very little system resistance to provide damping and a large peak develops. The decade method does better because the $100\mu\text{F}$ capacitors are nicely matched with the VRM output inductance but there are still some peaks above the target impedance. Component cost of large valued and sized capacitors is greater than the smaller values. The three methods are nearly equivalent in the chip/package resonance band of frequencies. The FDTIM capacitors are smaller on the average and use a higher percentage of low cost capacitors.

Conclusions

The Frequency Domain Target Impedance Method for selecting bypass capacitors has been discussed. The method uses quantities of several capacitor values to develop a flat impedance vs frequency profile. A major advantage is the damping provided for resonant peaks that may develop as capacitive or inductive components are added to the system. The method usually produces the lowest cost set of capacitors that will meet power quality specifications. There is however some additional complexity in manufacturing with several part number capacitors being assembled onto the PCB. The number of capacitors required to meet the target impedance up to a corner frequency is proportional to the corner frequency and inversely proportional to the target impedance. A formula has been given to estimate the number of PCB capacitors required.

References

- [1] L.D.Smith, R.E.Anderson, D.W.Forehand, T.J.Pelc, T.Roy, "Power Distribution System Design Methodology and Capacitor Selection for Modern CMOS Technology," IEEE Transactions on Advanced Packaging, Vol.22, No.3, August 1999, P284.
- [2] L.D.Smith, D.Hockanson, K.Kothari "A Transmission-Line Model for Ceramic Capacitors for CAD Tools based on Measured Parameters," Proc 52st Electronic Components & Technology Conference, San Diego, CA, May.2002, pp. 331-336.
- [3] L.D.Smith, "MLC Capacitor Parameters for Accurate Simulation Model," Design Con 2004.
- [4] I.Novak, Measuring Milliohms and PicoHenrys in Power Distribution Networks," Design Con 2000.
- [5] I.Novak, S.Pannala, J.R.Miller, "Overview of Some Options to Create Low-Q Controlled-ESR Bypass Capacitors," EPEP2004, October 25-27, 2004, Portland, OR.
- [6] L.D.Smith, R.Anderson, T.Roy, "Chip-Package Resonance in Core Power Supply Structures for a High Power Microprocessor," ASME Proceedings of Interpack'01, July,2001.
- [7] S.Weir, "Bypass Filter Design Considerations for Modern Digital Systems, A Comparative Evaluations of the Big "V", Multi-pole, and Many-pole Bypass Strategies," Design Con East 2005.