

Power Supply Design for PowerPC™ Processors

by *Gary Milliorn*
CPD Applications
Freescale Semiconductor, Inc.
Austin, TX

The internal logic of the Freescale Semiconductor host and networking processors, which contain a PowerPC™ core, typically consumes 90% or more of the total amount of device power, in comparison to other sections such as the I/O and PLL power. Consequently, the internal power, called core power or more simply V_{DD} , will place the heaviest demand on the power supply system.

This application note is intended to assist the system designer in delivering accurate and stable core power to the processor core logic.

1 Overview

The goal of the power supply section of any embedded system containing a PowerPC processor is to deliver a stable, accurate voltage to the processor. This is referred to as the power delivery system or PDS. A high-quality PDS is necessary to achieve the maximum operating frequency of the processor, and as current demand rises, the need to maintain a stable supply rises accordingly. Indeed, a common cause of system failure, or of an inability to operate a device at full speed, is often traceable to design errors in the PDS.

Contents

1. Overview	1
2. V_{DD} Power Supply Issues	4
3. Power Delivery System	7
4. Summary	18
5. References	19
6. Revision History	19

Overview

To accomplish the goal of achieving maximum performance and reliability, the power supply needs to be broken down into its individual pieces, and each designed carefully. The standard system model for representing the components of a typical system are shown in [Figure 1](#).

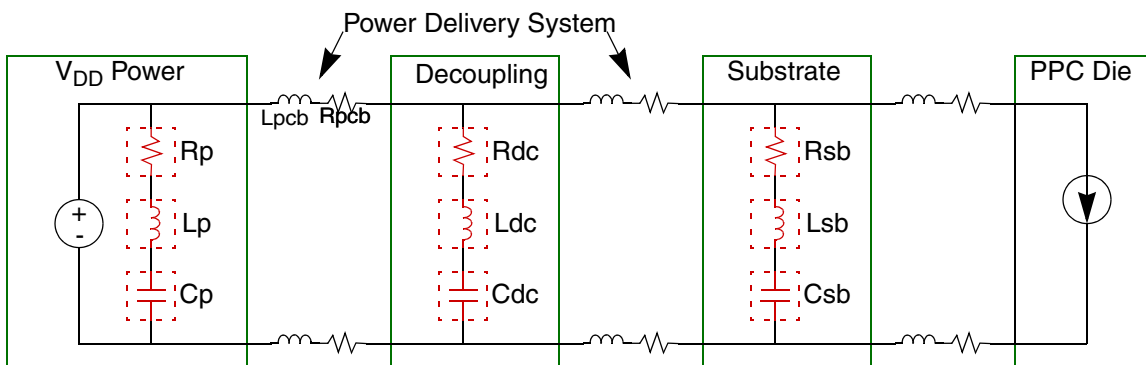


Figure 1. System Power Supply Block Diagram

The goal of this popular diagram is to remind the designer that these ‘real-world’ parasitics are omnipresent, and as current and frequency demands rise, are less likely to be ignorable.

The PDS can be divided into four categories:

- System power source (a linear or switching power supply)
- Power connections (planes/area fills, parasitics, cross-coupled noise, planar resonance, and so on)
- On-board decoupling
- The processor (the die and substrate considered as one)

Most of these devices are predetermined by the system architecture, leaving the power connections and the on-board decoupling as the means of addressing and preventing potential PDS problems.

1.1 The Role of Bypassing

Since the power supply is not a perfectly responsive DC source, when the processor current demand changes, the V_{DD} power drops unless bypassing components are present to ameliorate those effects. Figure 2 shows a typical transient event.

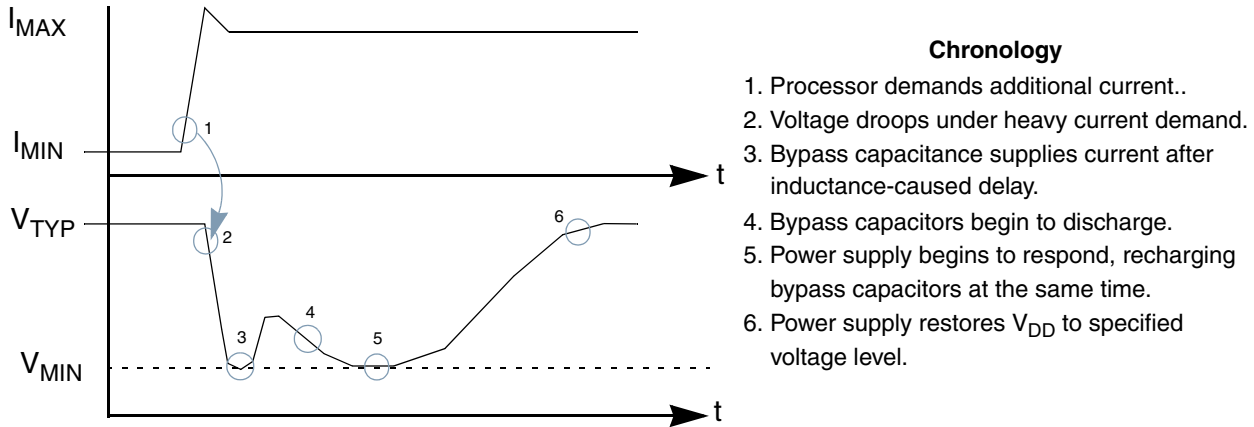


Figure 2. Power Supply Transient Event

Since the voltage level never falls below V_{MIN} , the PDS is working reliably. Although bypassing helps compensate for the droops, there are no perfect ‘magic’ bypass capacitors, so the goal is not to create a V_{DD} signal that appears to be DC (though that would certainly be nice), but instead to create a V_{DD} signal that is within the tolerance specifications of the processor. To do this with non-perfect devices, there needs to be an effective bypass system comprised of multiple, coordinated devices.

With both explicit device and implicit parasitic capacitance and inductances of various sizes distributed throughout the PDS, the frequency response curve of the system can be ‘lumped’ into distinct classes: low-, mid- and high-frequency bypass components. Figure 3 shows the bands.

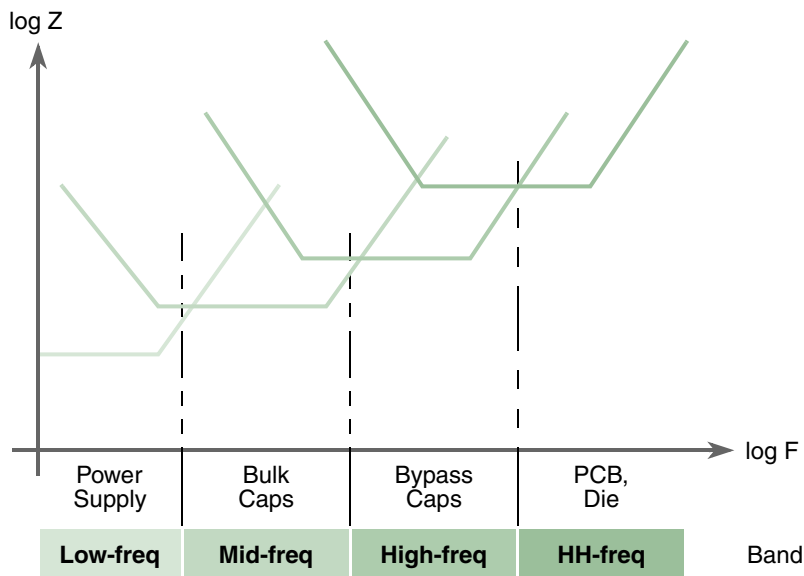


Figure 3. Bypass Frequency Bands

V_{DD} Power Supply Issues

(Though strictly speaking the highest-frequency resonance effects are affected by the PCB power plane, conventional naming is that the bypass capacitors are the high-frequency (HF) capacitors, so that term is used herein and the PCB/die plane are referred to as HHF.)

As the frequency changes, each part of the PDS responds proportionally; the low-impedance power supply responds to slow events, bulk capacitors to mid-frequency events, and so forth.

2 V_{DD} Power Supply Issues

This application note will not delve into the design of the power supply itself. Though it is assumed that a switching power supply is used (both for electrical and thermal efficiency), for both a switcher or a linear the output stage will be somewhat unique to each power supply, and the designer will need to follow guidelines presented in each power supply device datasheet.

Instead, treat the power supply and its output capacitance section as a black box. There are, however, several parameters specific to the processor, package and the printed circuit board which must be known to optimally design the power supply:

- Voltage supply value (or range)
- Voltage output accuracy (allowable ripple)
- Maximum power required
- Transient load change (di/dt)
- Target impedance of the PDS (Z_{TARGET})

Most of these are obtainable directly from the corresponding processor hardware datasheet. The parameter for di/dt is derived separately in [Section 2.1, “Measuring di/dt”](#).

Z_{TARGET} is defined as the PDS impedance needed to maintain V_{SUPPLY} at the rated I_{MAX} and V_{RIPPLE} as shown in the equation below. $VRIPPLEPCT$ is $VRIPPLE$ expressed as a percentage.

$$Z_{TARGET} = \frac{V_{SUPPLY} \times VRIPPLEPCT}{I_{MAX}}$$

Table 1 shows a summary of these parameters for a sampling of PowerPC processors at various speeds and voltages.

Table 1. PowerPC Processor Parameters

Processor	Parameter							No. of V _{DD} pins	Notes
	V _{SUPPLY}	P _{MAX}	I _{MAX}	V _{RIPPLE}	V _{RIPPLEPCT}	di/dt	Z _{TARGET}		
MPC8245 @ 300 MHz	1.8 V	2.2 W	1.22 A	± 100 mV	5.6%	0.1 A/ns	83 mΩ	20	1,2
MPC8245 @ 466 MHz	2.1 V	3.1 W	1.47 A	± 100 mV	4.8%	0.1 A/ns	69 mΩ	20	1
MPC7410 @ 500 MHz	1.8 V	11.9 W	6.61 A	± 100 mV	5.6%	0.2 A/ns	15 mΩ	18	1
MPC7445 @ 867 MHz	1.3 V	21.0 W	16.15 A	± 50 mV	3.9%	0.2 A/ns	3.1 mΩ	21	1
MPC7457 @ 1267 MHz	1.3 V	25.6 W	19.69 A	± 50 mV	3.9%	0.2 A/ns	2.6 mΩ	21	1,3
MPC7447A @ 1420 MHz	1.3 V	30.0 W	23.08 A	± 50 mV	3.9%	0.2 A/ns	2.2 mΩ	21	1

NOTES:

- 1 Subject to change without notice.
- 2 May be derated for the slower/cooler MPC8241 by scaling power through CV² F scaling.
- 3 See [9] in references for extensive details on di/dt measurement.

2.1 Measuring di/dt

The term di/dt obviously refers to the change in current over time. However, as shown in the system model in Figure 1, it also is a parameter of the silicon die that is essentially hidden by the capacitive and inductive components of the die substrate, the die-local bypass capacitors, the socket (if any) and other parasitics.

Consequently, the di/dt parameter used to design the power system is not the di/dt of the processor die (which would essentially be P_{MAX}/(2 × T_{CYC}); or 10A/ns for a 1GHz MPC7457), but the filtered di/dt of the combined processor, substrate-resident capacitors and the substrate itself. This di/dt is much slower, as the current demands are initially supplied by the adjacent transistors, die power traces, die substrate and local capacitors. As these devices are extremely close and/or tiny, the equivalent series inductance (ESL) itself is tiny and the effective frequency response range is very high (roughly the ‘HHF’ band in Figure 3). Even on a 1 GHz system, transients above 100–200 MHz should not be visible to the underlying PDS.

To measure the di/dt parameter for a particular system, the designer may model the PCB parameters (power planes, traces, and others) and capacitors (leads, via attach, and others) and simulate the environment.

Another way, perhaps the best way to measure di/dt is to directly attach a 50 Ω cable to a pair of adjacent BGA vias underneath the processor. Do not attach the wires to any nearby bypass capacitors, however convenient, else the measurement will include effects of the traces between the capacitor and the rest of the PDS. The BGA via attachment omits the higher frequency components of the power planes, but as previously discussed, this is accounted for already.

V_{DD} Power Supply Issues

Figure 4 shows a di/dt connection in place.




Figure 4. di/dt Cable Attachment

Figure 5 shows a resulting measurement.



Figure 5. di/dt Measurement

Note that this mirrors the response expected in [Figure 2](#). Performing repeated measurements over time and over a variety of circumstances allows the di/dt parameter to be graphically extracted from the data. This data is summarized in [9] of the references.

3 Power Delivery System

To design the PDS, first the kind and quantity of each class of bypass components must be determined. The second step is to determine the physical placement. Physical constraints may often limit the effectiveness of the bypass capacitors, and so occasionally component values and/or quantities will need to be adjusted for optimal results.

3.1 Bypass Capacitors

Selecting the optimal number, value and placement of the high-frequency bypass capacitors for the core power system may be a tricky problem. The issue is to design the system so the localized devices satisfy the mid-range transient response effectively, preventing droop and insuring a clean and stable power source.

The optimal solution is to completely model the components and power plane and study the simulated results, make changes to the placement and number of components. This is the only way to use only the bare minimum number required. Otherwise, use the calculations and provide additional coverage to compensate for unanticipated errors.

3.1.1 Bypass Capacitor Placement

The hardware specifications recommend one bypass capacitor per power pin (not including ground as a separate power). It can be difficult to place unless proper preparations are made beforehand. The first step is to use a BGA 'escape' pattern that radially expands from the center of the die. By dividing the BGA footprint into quadrants, a central axis is cleared down the vertical and horizontal portion of the PCB. The second step is to select small bypass capacitors (SMD 0603 or 0402 size). These may be placed in this central void area for very effective, localized bypassing. [Figure 6](#) shows just such a placement

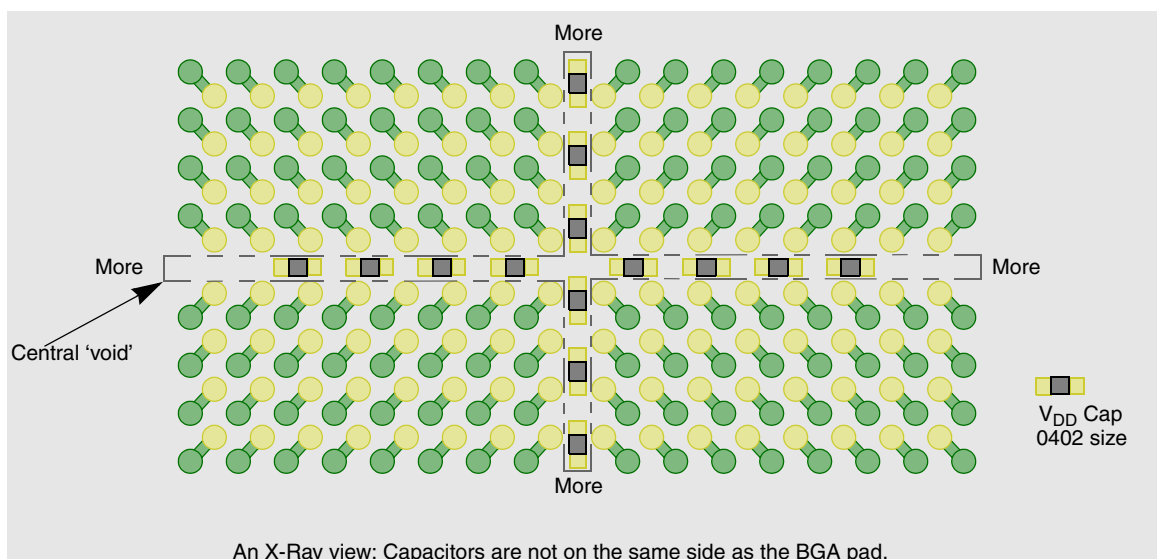


Figure 6. Centrally-Placed Bypass Capacitors

Power Delivery System

Once this area is opened up, approximately 50% of the 20 or so V_{DD} power bypass capacitors can be placed and directly attached to an adjacent power pad. However, not all will fit, so a second-order solution is needed for the remaining devices.

The key to placing the rest is to recognize that while placing bypass capacitors very far from a power pin reduces the effectiveness, capacitors are not required to be literally placed on the pin in order to have any effect whatsoever. As the trace length increases, or as the plane separation between the BGA pad and the capacitor increases, the inductance increases and the effectiveness of the capacitor decreases.

As shown in Figure 7, and described in the paper by Chen et.al., there is an effective radius, r_{eff} from the placement of a capacitor and where it remains useful [2].

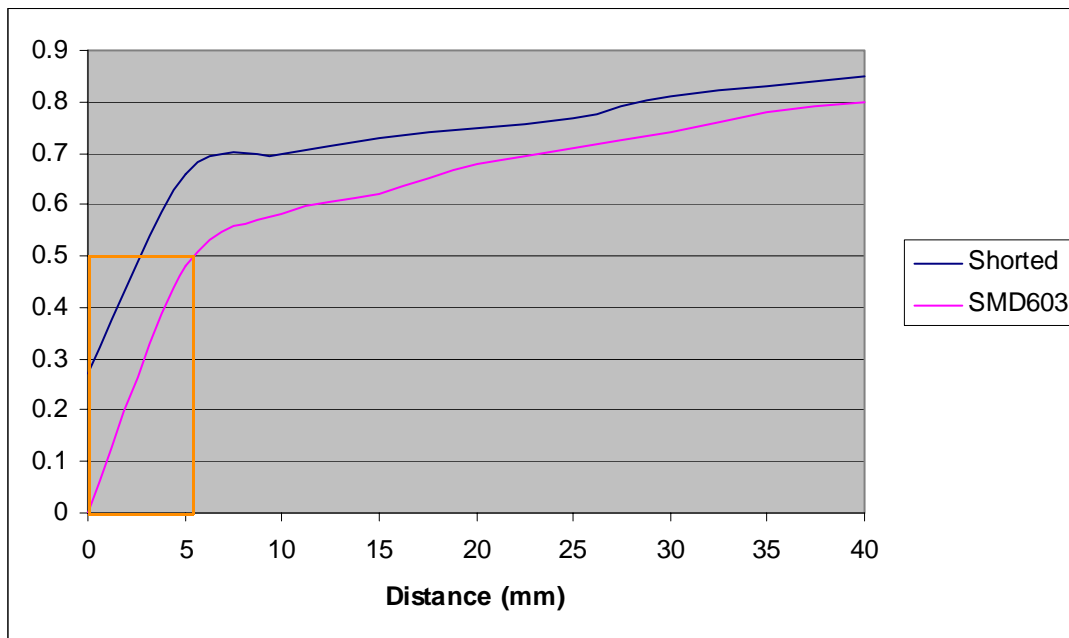


Figure 7. R_{EFF} for SMD 0603 Capacitors

As smaller devices become more common (for example, 0402, 0201), the r_{eff} will increase due to reduced inductance.

3.1.2 Bypass Capacitor Value

Selecting the bypass capacitor value is relatively easy: the package inductance is independent of the capacitance, and standard bypass values are well beyond their self-resonant frequencies (around 50 MHz) [10]. In such cases, the cut-off frequency is controlled by the inductance of the capacitors, and so they are chosen by the following criteria:

- Total number of capacitors
- Package of the capacitors
- Via attachment

Use capacitors 0508, 0402 or perhaps even a smaller size, with a value of 0.1 μF or greater.

NOTE

All capacitors may have the same value; there is no general benefit to be gained by deliberately adding a range of differing capacitance values (for example, using 0.1 mF, 0.01 mF, and 1 nF).

The rationale for using multiple capacitance values applied when capacitors had wire leads, where the inductance of the leads dominated parallel resonance equations. With small packages, the effect is so miniscule, any improvement could be equivalently achieved by halving the trace length of the power traces to the capacitor (approximately 6 dB for 0603 packages, less for 0402).

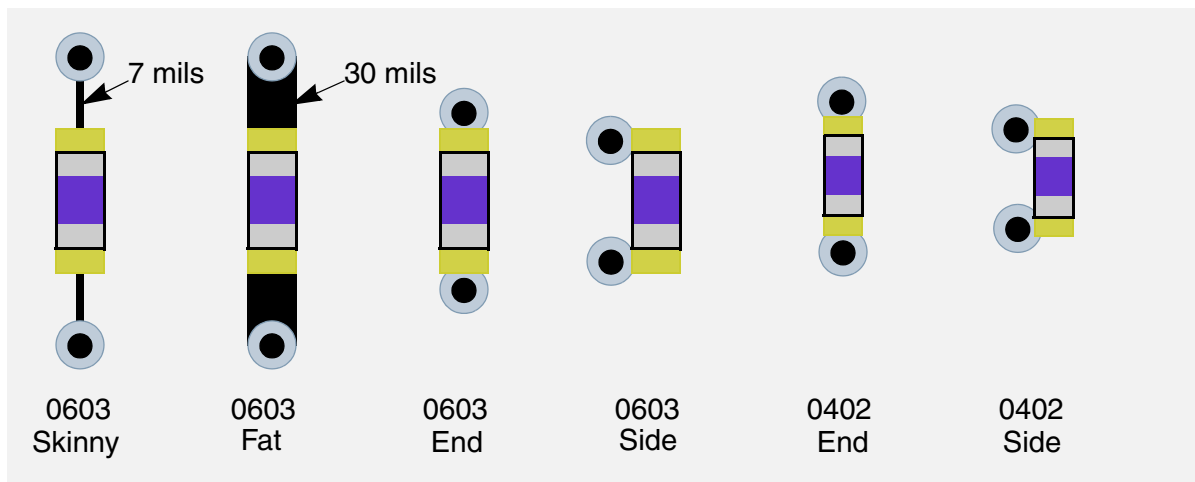
Note that this does not imply that there could not be a specific benefit to using a particular capacitance value; if there is a particularly noisy transient or radiated frequency, specific values may be of benefit. This is more of a filtering effect, and is typically an empirical correction which is made only after initial design is done. However, according to the previous section, replacing an existing 0.1 μF capacitor with a specific value will have no ill effect on the general bypassing design, as long as the capacitor is of the same type.

3.1.3 Bypass Capacitor Via Connections

Since the primary design goal of the bypass capacitors is to have a good, low-inductance connection, the quality of traces from the capacitor pad to their power/ground plane is critical. The total capacitor inductance is determined by the following:

- Capacitor inductance (which is determined by package)
- Pad/via inductance (which is determined by pad geometry and trace length from pad to via)
- PCB spreading inductance (which is determined by board (power plane) stackup)

Estimates are that 80% of the inductance is due to the capacitor itself, with the remaining 20% divided among the pad/via and PCB spreading inductance [6]. To minimize the latter (the former is covered in [Section 3.1.1, “Bypass Capacitor Placement”](#)), optimize the placement and routing of the bypass capacitor using the geometries shown in [Figure 8](#) and summarized in [Table 2](#).



See Johnson [8] for more details and dimensions.

Figure 8. Via Design for Bypass Capacitors

Table 2. Parasitic Inductance of Bypass Capacitors, nH


Hole Diameter 0.020 inches						
Via length: (inches)	0603 skinny	0603 fat	0603 end	0603 side	0402 end	0402 side
.004	1.51	0.89	0.42	0.33	0.38	0.21
.006	1.66	1.12	0.53	0.38	0.44	0.25
.010	2.13	1.47	0.68	0.51	0.58	0.32
.020	2.68	2.07	1.07	0.67	0.82	0.43
Hole Diameter 0.010 inches						
Via length: (inches)	0603 skinny	0603 fat	0603 end	0603 side	0402 end	0402 side
.004	1.51	0.95	0.50	0.36	0.42	0.26
.006	1.77	1.17	0.59	0.46	0.50	0.32
.010	2.18	1.52	0.77	0.61	0.67	0.40
.020	2.87	2.23	1.16	0.85	1.01	0.60

Note: Table from Johnson [8].

The goal for via placement is to minimize the inductance; at just under 4 nH/cm (~10 nH/in) just about any trace length to the pad can defeat the effect of the capacitor. Keep both traces to less than 10 mils total (0.010 in).

3.2 Bypass Capacitor Summary


Given the above restrictions, high-frequency bypass capacitors can be preferentially placed in the central void for core power, with the remaining capacitors placed as close as possible. [Figure 9](#) shows an example placement for the 360-pad BGA footprint, using centrally-placed capacitors and traces from the capacitor pads to the BGA pad.



An X-Ray view: Capacitors are not on the same side as the BGA pad.

Figure 9. 360 BGA Bypass Capacitor Arrangement

If via-in-pad is allowed, [Figure 10](#) is possible.



An X-Ray view: Capacitors are not on the same side as the BGA pad.

Figure 10. 360 BGA Bypass Capacitor Via-In-Pad Arrangement

For via-in-pad options, only one V_{DD} BGA ball does not have an adjacent ground pad, so it is connected with a short, fat trace. The central void area can be reused to provide OV_{DD} filtering, if needed.

3.3 Bulk Capacitors

Bulk capacitors are used to ‘refill’ the local core power supply and bypass capacitors during transient power events. The frequency response rate of bulk capacitors is much slower than the bypass capacitors, but much faster than the power supply (see Figure 3). Thus, while the requirement for low ESL devices remains, it is not quite as critical as it is for the bypass capacitors. In addition, the placement of the bulk capacitance is also not as critical, allowing them to be optimally placed simply by surrounding the device as shown in Figure 11, placing them as a group between the device and the power supply, or whatever (reasonable) alternative suits the design.




Figure 11. Peripherally-Placed Bulk Capacitors

3.3.1 Bulk Capacitor Value

Due to the low frequency response range of bulk capacitors, and the ESL of the device packages, the capacitance value and the ESR become the dominant selection criteria. The capacitance value is determined to compensate for the time the power supply takes to respond to transient power events presented to the system (note that this is *not* the same as the earlier-discussed di/dt value, it is the response caused by the di/dt of the processor).

The capacitance value is selected by the equation $C = I \times \frac{dt}{dV}$

where dV is the maximum allowable change in voltage, and dt is the response time of the power supply. For robust designs, it is essential to use the worst-case response time of the power source, not the typical or best-case.

For example, if a particular switcher can respond to a transient events within 200 ns, but may take as long as 15 μs under certain conditions (such as temperature or droop) then 15 μs is the correct value for dV. Thus, for a target such as the MPC7457 (from Table 1), the target must remain accurate to ± 50 mV, so C solves to $C = 20A \times \frac{15\mu s}{0.05V}$:

or 6000 μF. If the switcher is capable of responding faster, much less bulk capacitance will be required.

Note that this is irrespective of any requirements for localized output filtering capacitors of the power supply. Due to routing and other effects; output filtering capacitors are similar to bulk capacitors but serve different purposes.

Though values greater than 6000 μF exist, typically this capacitance value will be composed of several bulk capacitors. Low-ESR tantalums and organic electrolytic capacitors are generally suitable. When creating the bulk capacitor set, always select devices which help meet the required ESR value, as opposed to meeting the C value (that is, a higher C value is better than a higher ESR value).

Table 3 shows a sample of bulk capacitor parameters and is not intended as an approved vendor list since there are doubtless hundreds of other compatible devices. This is only a representative sampling of various devices and important selection criteria.

Table 3. A Sampling of Bulk Capacitor Parameters

Manufacturer	Type	Part Number	C	V _{Max}	ESR	Notes
AVX	Tantalum	TPME108K004R0018	1000 μ F	4 V	18 m Ω	Multi-anode for low ESR
AVX	Tantalum	TPSV108K004R0035	1000 μ F	4 V	35 m Ω	
Sanyo	PosCap	4TPE330MI	330 μ F	4 V	18 m Ω	
Sanyo	PosCap	2R5TPD1000M8	1000 μ F	4 V	8 m Ω	
Sanyo	OsCon	4SVPA680M	680 μ F	4 V	20 m Ω	
AVX	Tantalum	TAJY477K004R	470 μ F	4 V	900 m Ω	Don't use —ESR too high!

Continuing on with the previous design example, the required C value of 6000 μ F could be achieved with six AVX TPME108K004R0018 devices, but this provides an effective ESR of 3 m Ω , exceeding the Z_{TARGET} design goal. Adding an additional device, for a total of 7000 μ F, produces an ESR of 2.5 m Ω , meeting the required goal. To illustrate the importance of ESR for the bulk capacitance, consider instead the AVX TAJY477K004R capacitor shown in the table above. A perfectly fine device in some applications, but not for a power supply which may experience large transients. The high ESR value of the device means that 360 bulk capacitors would be needed!

3.3.2 Bulk Capacitor Via Connections

The vias of the bulk capacitors should generally follow the same optimized layout recommended for bypass capacitors (see Section 3.1.3, “Bypass Capacitor Via Connections”). Due to their generally larger size, placing vias under the device, or mechanically attaching them to the device pad is often easier.

In addition, insure that the via is rated for power use as opposed to signal vias; vias with larger holes have more plating surface area available to carry higher currents.

3.4 Plane Talk

Once the power supply is placed and the processor is adequately bypassed with high-frequency and bulk capacitance, all that remains is to connect the two via a power plane (or an area fill, or a planar connection). This seems simple enough, but careful considerations still remain.

Consider the power plane layout shown in [Figure 12](#).

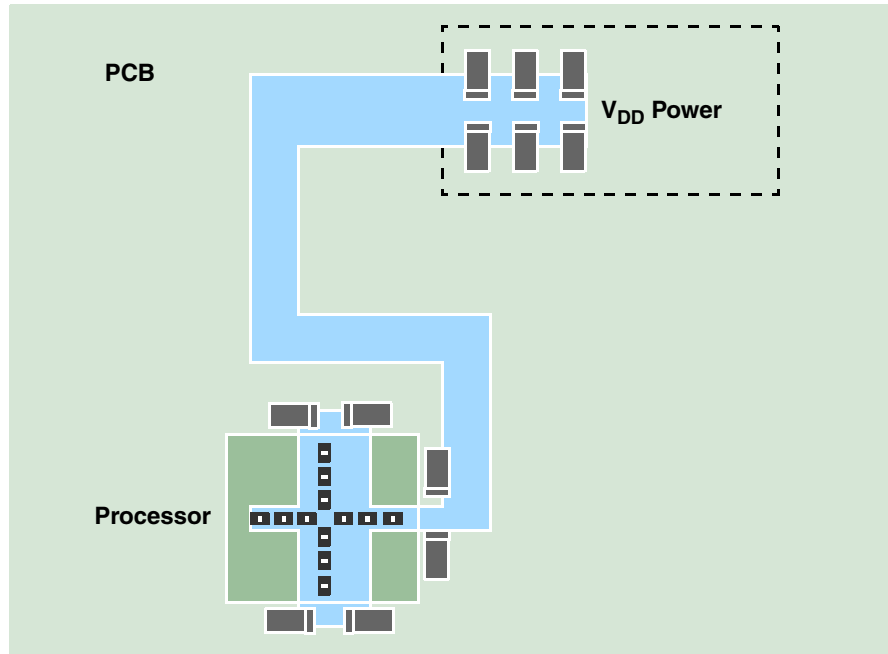


Figure 12. Power Plane Problems

Although at first the plane in [Figure 12](#) appears to be well designed, having followed the previous recommendations, a few problems are still present:

1. The power plane is long and circuitous, and will suffer from IR drops along the path.
2. The power plane pinches down to a narrow path and re-expands under the processor to connect to the processor V_{DD} pins.
3. The power plane is treated as if it were perfectly static.

The first issue can be addressed by preferentially assigning V_{DD} to a power plane in the PCB stackup. PCB layers are expensive, so the impulse to split power planes and share them can be enormous. However, poor planning at this stage can render a design useless which would be more expensive. The following list shows the preferred solutions starting with the best:

1. Dedicate an entire layer to V_{DD} . Do not share it with any other power.
2. Share V_{DD} with other power supplies other than OV_{DD} . Insure that the V_{DD} fill area is essentially a contiguous rectangle and is not impeded by numerous via voids.
3. Share V_{DD} and OV_{DD}/GV_{DD} , but preferentially treat V_{DD} fills as unchangeable, and make adjustments to OV_{DD}/GV_{DD} .

Note that with dedicated power plane pairs ($V_{DD} + GROUND$), there is a capacitor created that works at the highest frequencies as described in [Figure 3](#). As the capacitance is a function of total area (typically $\sim 100\text{pF}/\text{in}^2$), do not try to minimize the size of the V_{DD} plane, instead make it as large as possible. If that is not possible, ‘buried capacitance’ layers can be added between the two planes for even better effect.

Do not worry about signals victimizing the power plane; the planes are so much lower in impedance that signals cannot impose AC noise on them. Any transients on the power plane are due to current demands, not poor routing.

3.4.1 Area Fill Issues

The second issue noted earlier, the pinching down of an area-filled power plane, is easily resolved by not doing it in the first place. A solid plane suffers no such restrictions, as noted.

However, if it must be done, recall that a path is only as wide as its narrowest point. When flooding into the V_{DD} pin/via area under the processor, insure that the fill is of continuous width. In particular, avoid sharing V_{DD} and OV_{DD} on the same plane, as OV_{DD} and GV_{DD} pins more or less surround the V_{DD} pins on a 483-BGA package. This is particularly true if OV_{DD} and GV_{DD} attach to the same plane, as illustrated in Figure 13.

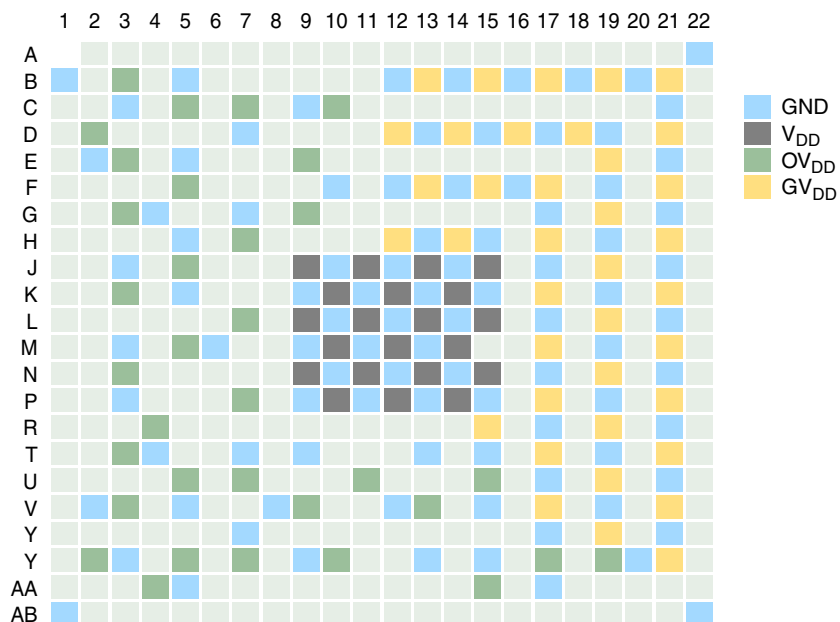


Figure 13. 483 BGA Power Pad Indication

In such a case, where shared planes are needed, optimize for the V_{DD} plane connection at all costs. In particular, the I/O power requirements are vastly lower than for the core, so attaching some of the OV_{DD} pads using traces will allow the V_{DD} plane to be routed out unscathed. An example is shown in Figure 14.

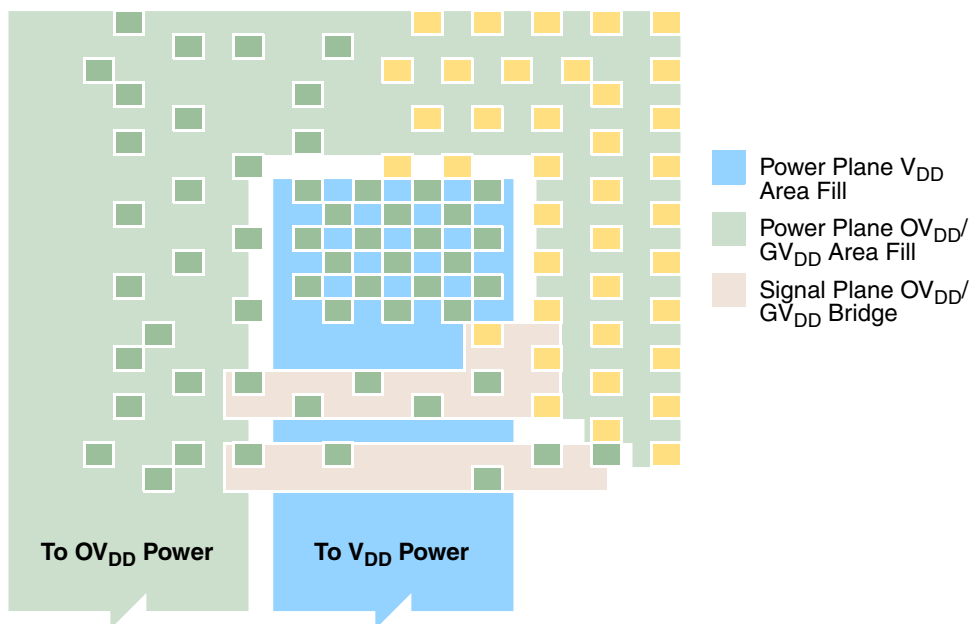


Figure 14. 483 BGA V_{DD} Surrounded Plane Escape

What this diagram shows is that the V_{DD} plane is escaped with no reduction in width, and attachments to the stranded OV_{DD} pads are made using a plane fill or fat traces on a signal layer. Note that this layout is still not optimal and still not preferred; in particular, it will be difficult to attach all the required bulk capacitance nearby.

3.4.2 Plane Resonance Issues

The third power plane design mistake is to regard the power plane as a perfectly static source. If the processor presented a perfectly stable current demand over all operating conditions (even during startup and shutdown), then the plane could perhaps be treated so simply. But as already seen in the discussion on di/dt , the power demands are dynamic and may fluctuate greatly.

When power demands are cyclic, as they definitely are from a high-speed processor making highly variable power demands on a cycle-by-cycle basis, the output of the power supply interacts with the imperfect lumped capacitance that is the power plane, and reflective signals can occur. Worse, those reflective signals can add to new transients, producing signals that are noise on top of noise. If this happens, the power plane is in a ‘resonant mode’ and may increase the likelihood the system will fail.

The easiest way to solve this is to load the system into a network simulator and look for resonances. However, if this is not practical or possible, use various rules of thumb to help forestall possible resonance modes. This requires visually evaluating the proposed power plane/area-fill and making necessary changes.

The rules of thumb (simulate if there are doubts) are summarized in [Table 4](#).

Table 4. Plane Resonance Protection ‘Rules of Thumb’

Capacitor	Value	Quantity	How Often
Bulk	Same as used for processor bulk.	1	Every 4 cm or when the plane changes directions.
Bypass	Same as used for processor bypass.	1	Every 1 cm.
Bypass	Same as used for processor bypass.	1	Every time the plane reduces in effective width (Which should be avoided).

3.4.3 Power Plane Redesigned

Using the material in the previous sections, the power plane can be redesigned as shown in [Figure 15](#).

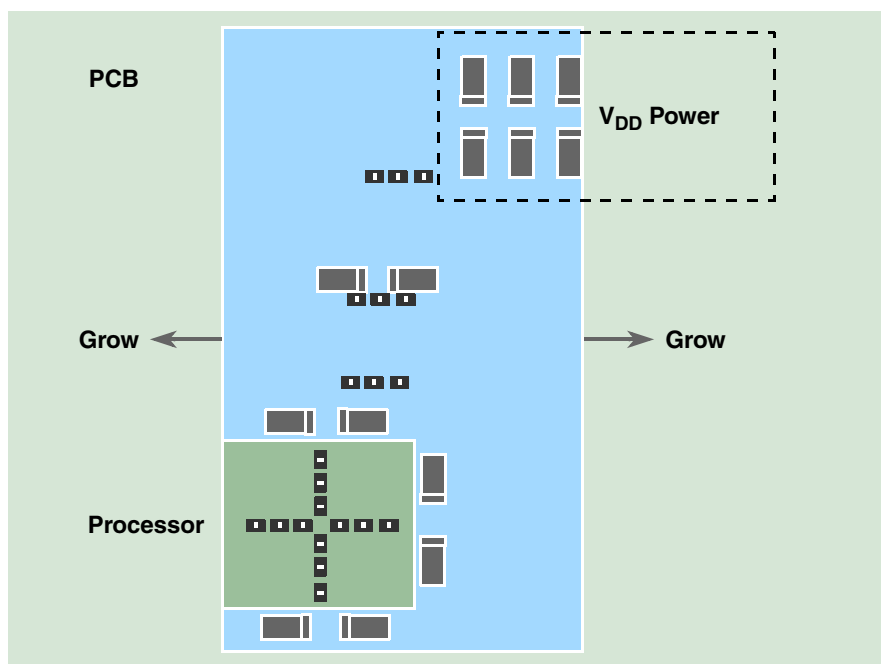


Figure 15. Plane Problems Eliminated—Best Case

In this best case, the power plane was opened up to fully encompass the processor. Bulk capacitance was added along the principal current path, at the proper interval, and additional bypass capacitors were added. The plane can be grown to the left and right as far as practical. Although it might seem somewhat ‘unfair’ to imply this is easily done (doubtless there are other components that might interfere), but having done the analysis of the plane requirements, it is best to start with this approach.

In short, it is the best practice to plan the PDS plane during placement, not just relegating the plane to whatever space can be ‘area filled’ after placement and routing are complete.

Summary

Alternately, if split planes are required, the layout shown in [Figure 16](#) might be used.

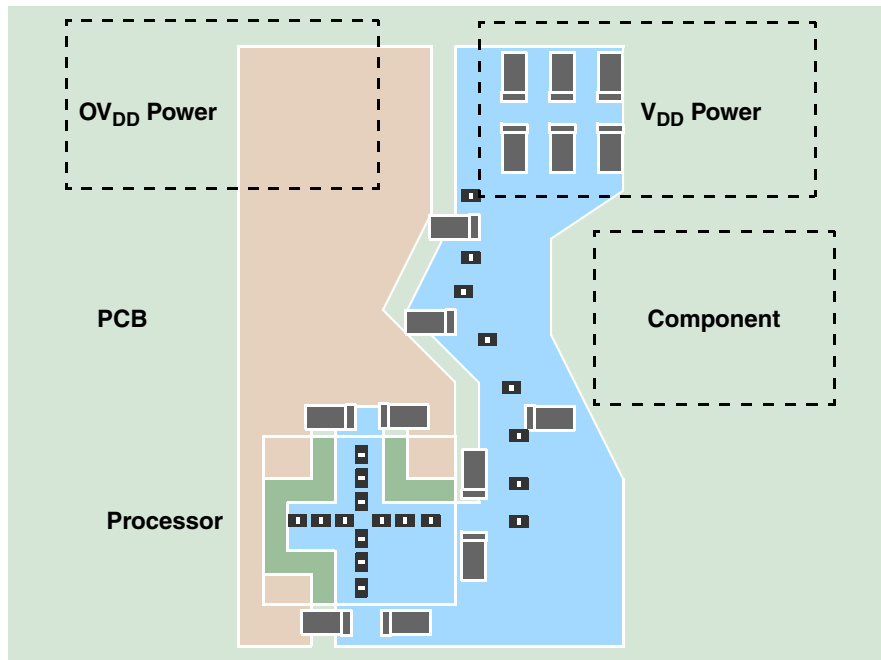


Figure 16. Plane Problems Eliminated—Next Best Case

In this case, the V_{DD} plane has been restricted to allow routing an OV_{DD} plane/area-fill and avoiding another component. Additional capacitors were added to insure self-resonance does not occur, but the plane is uniformly wide and is expanded into unused areas to increase the capacitance and HHF response.

Note that, although not shown, the OV_{DD} plane would benefit from the same consideration given the V_{DD} plane (such as distributed capacitance, growth, and so forth).

4 Summary

[Table 5](#) summarizes the rules developed in this application note

Table 5. Summary of Power System Design Rules

Importance	Rule
1	Use a solid V_{DD} plane or split wisely.
2	Plan the power plane during placement.
3	Place high-frequency bypass capacitors directly under the processor.
4	Use smallest package/lowest ESL high-frequency bypass capacitors.
5	Use the shortest trace length between the HF cap via and the pad.
6	Use lowest ESR bulk capacitors.
7	Make the V_{DD}/OV_{DD} plane as large as possible.
8	Place remaining high-frequency bypass capacitors around the processor.
9	Place bulk capacitors surrounding the processor.

By following these guidelines, the power delivery system will be robust and will allow the device to operate over its fullest frequency, temperature, and voltage ratings.

5 References

The following resources were used in the creation of this application note. Interested readers should refer to these publications for further details. In particular, Dr. Johnson's book *High Speed Digital Design: A Handbook of Black Magic* deserves a place on every engineer's desk.

- [1] Clayton R. Paul, "Effectiveness of Multiple Decoupling Capacitors," *IEEE Transactions on Electromagnetic Compatibility*, Vol. 34, No. 2, pp. 130-133, May 1992.
- [2] Huabo Chen, Jiayuan Fang, and Weimin Shi, "Effective Decoupling Radius of Decoupling Capacitor," *IEEE Electrical Performance of Electronic Packaging*, pp. 277-280, Oct. 2001.
- [3] Yun Chase, "Introduction to Choosing MLC Capacitors For Bypass/Decoupling Applications," AVX Corporation, <http://www.avx.com/docs/techinfo/mlcbypas.pdf>.
- [4] Feng Yu, "Power Integrity Analysis in a split plane PCB", Ansoft, ftp://ftp.ansoft.com/techsup/download/web/articles/Zhejiang_Univ.pdf (Oct. 2002).
- [5] Howard W. Johnson and Martin Graham, *High Speed Digital Design: A Handbook of Black Magic*, Prentice Hall PTR, 1993.
- [6] Ray Anderson, <http://www.qsl.net/wb6tpu/si-list2/0223.html>.
- [7] Howard Johnson, "Parasitic Inductance of a Bypass Capacitor," <http://signalintegrity.com/Pubs/edn/ParasiticInductance.htm>.
- [8] Howard Johnson, "Parasitic Inductance of a Bypass Capacitor II," http://signalintegrity.com/Pubs/news/6_09.htm.
- [9] Malcolm Ward, "Power Consumption and Power Supply Design Requirements for the Latest PowerPC Host Processors," SNDF2004 Presentation, https://www.freescale.com/webapp/Download?colCode=SNDF2004_DALLAS_H1105&nodeId=02VS011Cc5pzMP551361586320&appType=license&location=overview.
- [10] Steve Weir, <http://www.freelists.org/archives/si-list/02-2004/msg00259.html>.

6 Revision History

Table 6 provides a revision history for this application note.

Table 6. Document Revision History

Rev. No.	Date	Substantive Change(s)
1	08/23/04	Initial release.
1.1	09/02/04	Minor editing

How to Reach Us:

USA/Europe/Locations Not Listed:

Freescale Semiconductor
Literature Distribution Center
P.O. Box 5405,
Denver, Colorado 80217
1-480-768-2130
(800) 521-6274

Japan:

Freescale Semiconductor Japan Ltd.
Technical Information Center
3-20-1, Minami-Azabu, Minato-ku
Tokyo 106-8573, Japan
81-3-3440-3569

Asia/Pacific:

Freescale Semiconductor Hong Kong Ltd.
2 Dai King Street
Tai Po Industrial Estate
Tai Po, N.T. Hong Kong
852-26668334

Home Page:

www.freescale.com

Information in this document is provided solely to enable system and software implementers to use Freescale Semiconductor products. There are no express or implied copyright licenses granted hereunder to design or fabricate any integrated circuits or integrated circuits based on the information in this document.

Freescale Semiconductor reserves the right to make changes without further notice to any products herein. Freescale Semiconductor makes no warranty, representation or guarantee regarding the suitability of its products for any particular purpose, nor does Freescale Semiconductor assume any liability arising out of the application or use of any product or circuit, and specifically disclaims any and all liability, including without limitation consequential or incidental damages. "Typical" parameters which may be provided in Freescale Semiconductor data sheets and/or specifications can and do vary in different applications and actual performance may vary over time. All operating parameters, including "Typicals" must be validated for each customer application by customer's technical experts. Freescale Semiconductor does not convey any license under its patent rights nor the rights of others. Freescale Semiconductor products are not designed, intended, or authorized for use as components in systems intended for surgical implant into the body, or other applications intended to support or sustain life, or for any other application in which the failure of the Freescale Semiconductor product could create a situation where personal injury or death may occur. Should Buyer purchase or use Freescale Semiconductor products for any such unintended or unauthorized application, Buyer shall indemnify and hold Freescale Semiconductor and its officers, employees, subsidiaries, affiliates, and distributors harmless against all claims, costs, damages, and expenses, and reasonable attorney fees arising out of, directly or indirectly, any claim of personal injury or death associated with such unintended or unauthorized use, even if such claim alleges that Freescale Semiconductor was negligent regarding the design or manufacture of the part

Learn More: For more information about Freescale Semiconductor products, please visit www.freescale.com

Freescale™ and the Freescale logo are trademarks of Freescale Semiconductor, Inc. The PowerPC name is a trademark of IBM Corp. and is used under license. All other product or service names are the property of their respective owners.

© Freescale Semiconductor, Inc. 2004.

AN2747
Rev. 1.1
09/2004