



AMD Athlon™

Processor

Voltage Regulation Design

Application Note

© 2000 Advanced Micro Devices, Inc. All rights reserved.

The contents of this document are provided in connection with Advanced Micro Devices, Inc. ("AMD") products. AMD makes no representations or warranties with respect to the accuracy or completeness of the contents of this publication and reserves the right to make changes to specifications and product descriptions at any time without notice. No license, whether express, implied, arising by estoppel or otherwise, to any intellectual property rights is granted by this publication. Except as set forth in AMD's Standard Terms and Conditions of Sale, AMD assumes no liability whatsoever, and disclaims any express or implied warranty, relating to its products including, but not limited to, the implied warranty of merchantability, fitness for a particular purpose, or infringement of any intellectual property right.

AMD's products are not designed, intended, authorized or warranted for use as components in systems intended for surgical implant into the body, or in other applications intended to support or sustain life, or in any other application in which the failure of AMD's product could create a situation where personal injury, death, or severe property or environmental damage may occur. AMD reserves the right to discontinue or make changes to its products at any time without notice.

Trademarks

AMD, the AMD logo, AMD Athlon, and combinations thereof are trademarks of Advanced Micro Devices, Inc.

Microsoft, Windows, and Windows NT are registered trademarks of Microsoft Corporation.

Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

Contents

Revision History	ix
Introduction	1
Processor Power Requirement	2
Voltage Plane	2
Power Supply Specification	2
Electrical Specifications	2
Voltage Identification (VID) Settings	3
Selecting a Power Supply Design	4
Switching Regulator Design	5
Physical Requirements	7
Decoupling Recommendations	10
Power Distribution	10
Current Transient Response	11
Examples	11
Output Voltage Response Measurement Techniques	13
Output Voltage Response Measurement Utilities	14

List of Figures

Figure 1.	Linear and Switching Voltage Regulators	4
Figure 2.	Basic Asynchronous Design	5
Figure 3.	Basic Synchronous Design	6
Figure 4.	Power Supply Area Allotment	8
Figure 5.	Power Distribution Model	10
Figure 6.	Bulk Decoupling versus Output Voltage Response for 1.6 V @ 35 A	13

List of Tables

Table 1.	Voltage and Current Requirements	2
Table 2.	Voltage Identification (VID) Codes	3
Table 3.	Physical Dimensions of Power Supply	8

Revision History

Date	Rev	Description
February 2000	E-1	Added the 850-MHz AMD Athlon™ processor to Table 1, “Voltage and Current Requirements,” on page 2.
January 2000	E	Added V_{CC} information for the 800-MHz AMD Athlon™ processor to Table 1, “Voltage and Current Requirements,” on page 2.
October 1999	D	Revised Table 1, “Voltage and Current Requirements,” on page 2—Changed the min V_{CC} transient tolerance to -50mV and Eff_{min} to 70%.
October 1999	C	Changed the core regulator current specification to 35A throughout the document.
August 1999	B	Initial public release.

Application Note

AMD Athlon™ Processor Voltage Regulation Design

Introduction

The AMD Athlon™ processor is a high-performance x86-compatible processor with over 22 million transistors. Due to the large number of transistors that can switch simultaneously, power supply designs must meet large transient power requirements.

This application note is intended to guide the board designer through the process of developing a reliable power supply that meets the low-voltage, high-current demands of the AMD Athlon processor. The goal is to design a solution that works over a wide voltage range and a 3-amp (A) to 35-A current range. This large range allows motherboard designers to prepare for the next generation of processors from AMD. In addition, this application note provides basic guidelines on circuit decoupling for reduction of noise generated by fast current transients.

AMD encourages designers to provide flexibility to support multiple voltages in their designs. To allow multiple core voltages for the processor, support of the voltage identification (VID) codes shown in Table 2 on page 3 is required. By providing flexibility in the power design now, lower voltage parts available in the future may require little or no change to the motherboard.

Processor Power Requirement

Voltage Plane

One supply voltage, V_{CC} , is required to support the processor.

The power-supply finger-edge assignments for the Slot A cartridge are as follows:

$V_{CC_Core}[1 - 44]$: A24, A82, A52, A38, A86, A98, A58, A71, A30, A32, A104, A76, A64, A62, A44, A110, A60, A26, A48, A92, A90, A84, A80, A36, A40, A78, A100, A56, A50, A54, A68, A66, A42, A34, A102, A96, A46, A88, A94, A22, A74, A107, A28

For a detailed listing of signal names and signal locations on the module, see the *AMD Athlon™ Processor Data Sheet*, order# 21016.

Power Supply Specification

Electrical Specifications

Voltage and current specifications are provided in the *AMD Athlon™ Processor Data Sheet*, order# 21016. However, additional information is required to design an appropriate power supply for the AMD Athlon processor. Table 1 shows the electrical requirements of the processor not included in the data sheet along with power supply design targets. These projected values allow designers to support future processors on the same motherboard, thereby allowing upgradeability.

Table 1. Voltage and Current Requirements

	Description	Min	Typ	Max
V_{CC}	V_{CC} for the 800-MHz and 850-MHz AMD Athlon™ processors		1.70 V	
	V_{CC} for the 500-MHz through 750-MHz AMD Athlon processors		1.60 V	
	V_{CC} static tolerance at the regulator	-0.050 V		+0.050 V
	V_{CC} transient tolerance at the regulator	-0.050 V		+0.100 V
I_{CC}	I_{CC} for the processor—lower current	3 A		35 A
dI_{CC}/dt	Current slew rate			30 A/ μ s
Eff_{min}	Efficiency of the regulator at minimum load	70%		
Eff_{max}	Efficiency of the regulator at maximum load	80%		

Voltage Identification (VID) Settings

As the trend toward smaller process geometries continues (0.25-micron to 0.18-micron), the processor core voltage is projected to drop. To provide maximum flexibility for upgrading a motherboard, switching regulator controllers with the 5-bit VID code are required. Using this feature, processors still in the design stage can be supported, as long as they do not exceed the current limit of the design.

The AMD Athlon processor has four pins that indicate the voltage requirements of the device—VID[3], VID[2], VID[1], and VID[0]. The arrangement of these four signals sets the voltage of the processor. Industry-standard VID settings also include VID[4]. However, on the AMD Athlon processor, VID[4] is set low (0b) internally on the processor module. The VID values that must be supported for the AMD Athlon are shown in Table 2.

Table 2. Voltage Identification (VID) Codes

Module Signal Name					Core Voltage of Processor
VID[4]	VID[3]	VID[2]	VID[1]	VID[0]	
0	1	1	1	1	1.30 V
0	1	1	1	0	1.35 V
0	1	1	0	1	1.40 V
0	1	1	0	0	1.45 V
0	1	0	1	1	1.50 V
0	1	0	1	0	1.55 V
0	1	0	0	1	1.60 V
0	1	0	0	0	1.65 V
0	0	1	1	1	1.70 V
0	0	1	1	0	1.75 V
0	0	1	0	1	1.80 V
0	0	1	0	0	1.85 V
0	0	0	1	1	1.90 V
0	0	0	1	0	1.95 V
0	0	0	0	1	2.00 V
0	0	0	0	0	2.05 V

Notes:
0 = Low, GND
1 = High, V_{CC}

Selecting a Power Supply Design

Platforms today require DC-to-DC voltage conversion circuits to supply lower voltages to the processor core and I/O. Two types of regulators are used—linear and switching.

A linear regulator provides excellent dynamic-load response in the low-voltage, high-current environment. It also contributes to simplified design and lower cost. However, the efficiency loss and heat generated by a linear regulator is high for the current levels targeted. As processor voltages drop and currents increase, it becomes more difficult to implement a linear solution due to increased inefficiencies. Therefore, due to the implications linear solutions create, a switching regulator should be considered for the AMD Athlon processor.

A switching regulator meets the efficiency and size limitations of low-cost, high-performance desktop designs. Figure 1 shows a simplified example of a linear and switching regulator. The switching regulator uses a series switch in conjunction with the output capacitor (C_{OUT}) to control the ON/OFF ratio in order to obtain an average output voltage. Because the switch frequently turns off, only a small amount of power is lost during conversion.

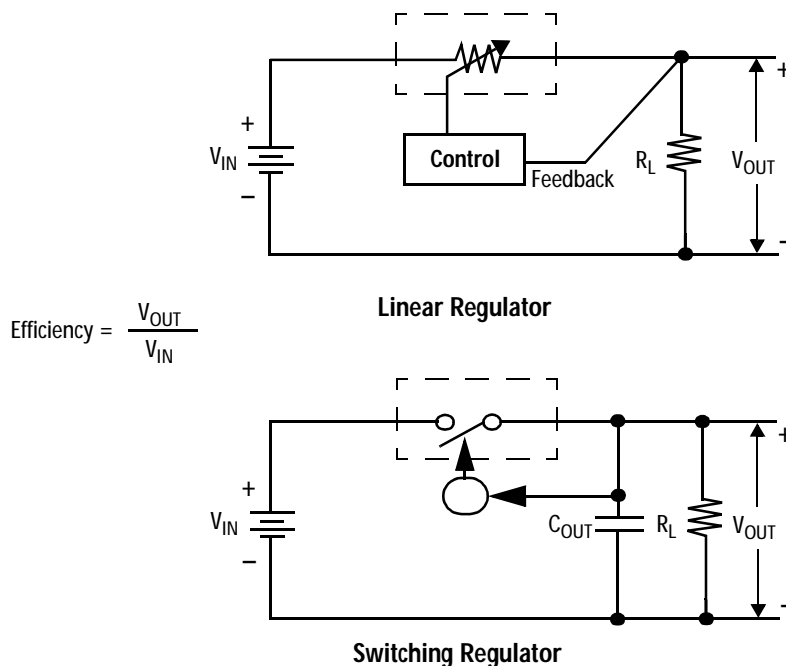


Figure 1. Linear and Switching Voltage Regulators

Switching Regulator Design

A switching regulator varies the switch duty cycle (ON/OFF ratio) according to the output feedback. A large output capacitor (C_{OUT}) is used in the switching design to achieve a constant average output. The switching regulator delivers higher efficiency than a linear regulator, but the tradeoffs are higher ripple voltages (noise) and slower transient current response time. A series inductor is used to supply current to the load during the switch OFF time, adding complexity to the design. In addition, the inductor and the output capacitor increase the overall cost of the switching regulator design, relative to a linear regulator design.

The power supply design must account for a low current (I_{CC}) drain when the processor enters the Stop Grant state. The power supply must ensure that the minimal current drain does not cause any adverse side effects (drift out of regulation, over-compensation, or shutdown) that could corrupt or damage the functionality of the processor.

In the basic asynchronous circuit design shown in Figure 2, Q1 turns on to charge C_{OUT} and builds up the magnetic field in L1. When the feedback from the sense input is too high, the controller turns off Q1. Current is supplied to the load by the collapsing magnetic field in L1 and the discharge of C_{OUT} . When the sense feedback detects a drop in the load voltage, the controller turns on Q1 to recharge the circuit. CR2 supplies a return path for L1 when it is supplying current. This design is less efficient than a synchronous design primarily because the power dissipated in CR2 is higher than Q2 in the synchronous design.

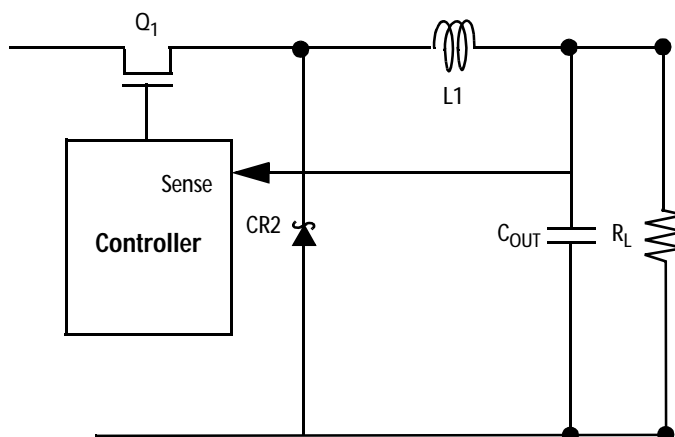


Figure 2. Basic Asynchronous Design

The operation of the basic synchronous circuit design shown in Figure 3 is essentially the same as the asynchronous design. Q1 turns on to charge C_{OUT} and builds up the magnetic field in L1. When the feedback from the sense input is too high, the controller turns Q1 off. Current is supplied to the load by the collapsing magnetic field in L1 and the discharge of C_{OUT} . When the sense feedback detects a drop in the load voltage, the controller turns on Q1 to recharge the circuit. Q2 supplies a return path for L1 when it is supplying current. When Q1 is on, Q2 is off and when Q1 is off, Q2 is on. The main reason this design is more efficient is because the power dissipated in Q2 is lower than the power in CR2.




Figure 3. Basic Synchronous Design

Another consideration is power dissipation in the lower MOSFET (synchronous) or diode (asynchronous). As the output voltage decreases, the power dissipation in CR2 (Q2) increases. The higher power dissipation may require using a different package type or adding a heatsink to dissipate the additional power.

To determine if the transistors or the diode need a heatsink, use the following equation:

$$P = I^2R \cdot \text{duty cycle (Q1)}$$

$$P = I^2R \cdot (1 - \text{duty cycle (Q2)})$$

$$\text{Duty cycle} \sim V_{OUT}/V_{IN}$$

Compare these calculations with the specifications of the device used.

The simple examples shown in Figure 2 and Figure 3 are single-phase designs. *Single phase* is a term used to describe the number of drive FET banks that turn on for a given switching cycle. In Figure 3, Q1 is the only drive FET supplying current to the processor. Therefore, this design would be a single-phase design.

A *multi-phase* design incorporates more FETs, with each FET switching on and off at different times. Each bank of FETs requires its own inductor, capacitors, and lower drive discrete components (diode for asynchronous, FET for synchronous, or a combination of diode and FET). However, smaller discrete components can be used because each bank divides the total current by the number of banks available. For example, a two-phase switching regulator design has two banks, each capable of handling one half the total current. Likewise, each bank in a four-phase switching regulator is capable of handling one fourth the total current.

To assist in meeting the demanding efficiency requirements detailed in Table 1 on page 2, several key design decisions must be made to obtain higher efficiency. Each design choice has a cost/performance trade-off. First, one of the key motherboard components is the power transistor. The transistor can be replaced with one that has a lower $R_{DS(ON)}$ (resistance-drain-to-source when the transistor is on) or two transistors can be paralleled. Of course, higher performing transistors usually carry a higher cost, but performance is greatly improved. Secondly, another design trade-off to obtain higher efficiency is designing a synchronous switching regulator versus a non-synchronous switching regulator. As previously mentioned, the power loss in the lower drive is substantially less in a FET versus a diode.

Physical Requirements

The cost of switching regulator bill of materials (BOM) increases as the current capability of the design increases. Due to the large current required to operate the processor, a single-phase regulator would occupy a large surface area on the motherboard. Because an inductor carrying 30A is physically larger than an inductor carrying 10A, the layout must allow sufficient space. Numerous large capacitors and a large inductor would occupy a very large percentage of the motherboard at the targeted currents.

A designer should consider a multi-phase design for the reduced component size it requires. A multi-phase design would potentially have a lower cost BOM when compared to a single-phase design at the targeted current levels.

A physical constraint is placed on the area allowed for a power supply. Design of a power supply must take into consideration the area allotment detailed in Table 3. By designing to these size constraints, a motherboard designer is more likely to choose the smaller, lower cost design that requires minimal board layout effort.

Table 3. Physical Dimensions of Power Supply

Symbol	Description	Maximum
L	Length	5.0 inches
W	Width	1.0 inches
H	Height	2.5 inches

The entire power solution, including FETs, inductors, and capacitors, must fit in an area on the motherboard of approximately five square inches. Figure 4 shows the board area allotted on motherboards to house the power supply solution for the processor.




Figure 4. Power Supply Area Allotment

More size flexibility is available if the motherboard designer budgets more area for the power regulator. The motherboard designer and the power regulator supplier can decide to provide a solution that meets and/or exceeds the specifications detailed in this document. As long as these specifications are met, the power regulation solution can occupy as much board area as necessary.

Decoupling Recommendations

Power Distribution

In order to maintain a stable voltage supply during fast transients, power planes with high-frequency and bulk decoupling capacitors are required. Figure 5 shows a power distribution model for the power supply and the processor. The bulk capacitors (C_B) are used to minimize ringing, and the processor decoupling capacitors (C_F) are spread evenly across the circuit to maintain stable power distribution.

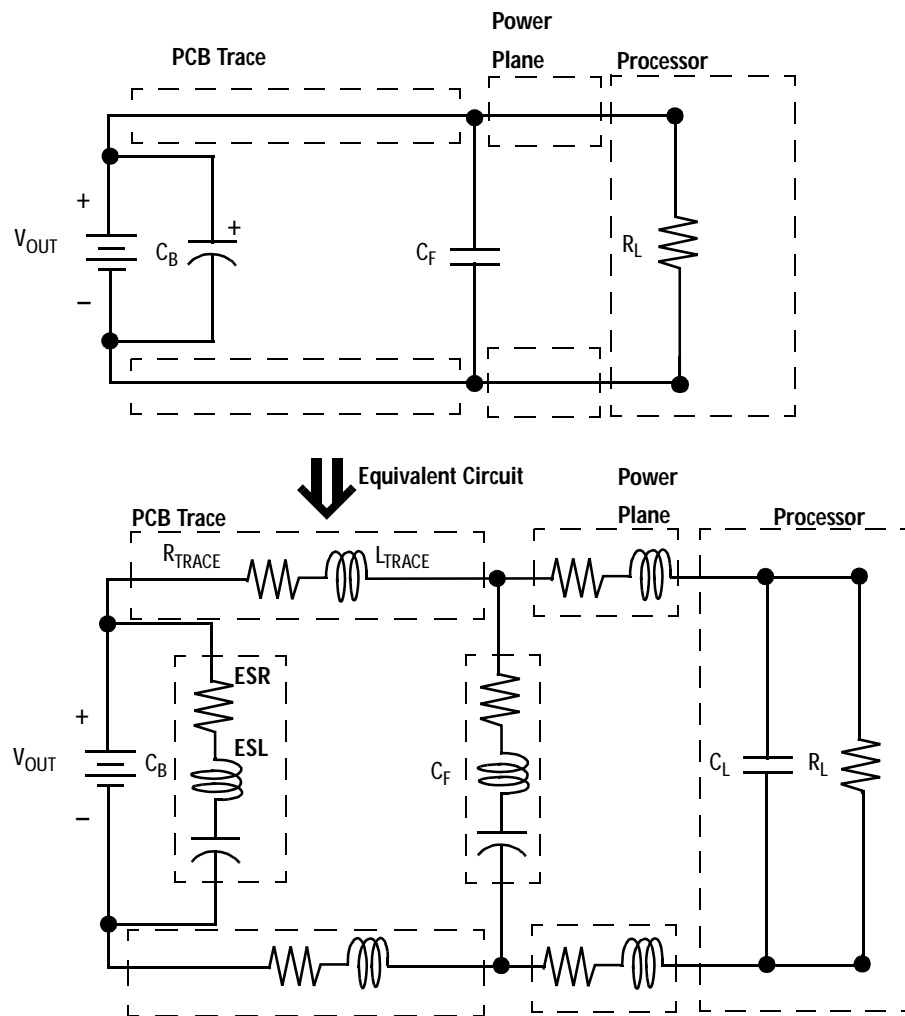


Figure 5. Power Distribution Model

Current Transient Response

In the power distribution model shown in Figure 5, C_B represents bulk capacitors for the power supply and C_F represents high-frequency capacitors for processor decoupling. The C_F capacitors are designed into the AMD Athlon processor and have a total value of 125.8 μ F. The bulk capacitors supply current to the processor during sudden excessive current demands that cannot be supplied by the voltage regulator (for example, a transition from the Stop Grant state to normal mode). The required C_B is calculated using the following equation (ideal case):

$$C_B \geq \frac{\Delta I}{\Delta V} \cdot \Delta t$$

Where:

- ΔI is the maximum processor current transient
- ΔV is the tolerance times the nominal processor voltage
- Δt is the voltage regulator response time

Examples

The following examples are not the only solutions. Based on the availability of parts and the choice of controller, many correct solutions are possible. The examples, which use multi-layer ceramic (MLC) and Oscon capacitors, are intended to give insight into the requirements and not to specify a particular solution. The use of tantalum and aluminum electrolytic capacitors is acceptable as long as good quality, low equivalent series resistance (ESR) parts are used.

Example 1: Theoretical 1.6V @ 30A

Assuming the maximum processor current transient is 30A, the voltage tolerance of the processor is less than 100mV (as specified in Table 1 on page 2) and the voltage regulator response time is 2 μ s, the minimum capacitance for the bulk decoupling is:

$$C_B \geq (30A/0.100V) \cdot 2\mu s = 600\mu F$$

ESR and ESL (equivalent series inductance) are introduced in the model shown in Figure 5 on page 10. C_B contains ESR and ESL, causing a voltage drop during current transient activity. The resistive and inductive effect of the capacitors must be

taken into account when designing processor decoupling. Low-ESL and low-ESR capacitors should be used to obtain better voltage and current output characteristics. Taking into account the ESR, the following equation is used to calculate C_B :

$$C_B \geq \frac{\Delta I}{(\Delta V - (\Delta I \cdot \text{ESR}))} \cdot \Delta t$$

Example 2:
Actual 1.6V @ 30A

This example assumes the maximum processor current transient is 30A, the voltage tolerance of the processor is less than 100 mV ($1.6\text{ V} \pm 100\text{ mV}$), and the voltage regulator response time is $2\mu\text{s}$.

With 10 MLC capacitors ($22\mu\text{F}$) with $15\text{-m}\Omega$ ESR each and one Oscon capacitor ($820\mu\text{F}$) with $12\text{-m}\Omega$ ESR (the parallel resistance is $1.33\text{m}\Omega$) used as bulk capacitors, the minimum bulk capacitance is:

$$C_{OUT} \geq ((30\text{A}/(0.100\text{V} - [30\text{A} \cdot 1.33\text{m}\Omega])) \cdot 2\mu\text{s} = 1000\mu\text{F}$$

Example 3:
Targeted 1.6V @ 35A

This example assumes a device with a maximum processor current transient of 35A, a processor voltage tolerance of less than 100 mV, and a voltage regulator response time of $2\mu\text{s}$.

Using eighteen MLC capacitors ($22\mu\text{F}$) with $15\text{-m}\Omega$ ESR each and one Oscon capacitor ($820\mu\text{F}$) with $12\text{-m}\Omega$ ESR (the parallel resistance is $0.526\text{m}\Omega$) as bulk capacitors, the minimum bulk capacitance is:

$$C_{OUT} \geq ((35\text{A}/(0.100\text{V} - [35\text{A} \cdot 0.779\text{m}\Omega])) \cdot 2\mu\text{s} = 962.40\mu\text{F}$$

Figure 6 on page 13 shows the voltage curve for this case.

By following this example, a designer allows for motherboard upgradeability using the AMD Athlon processor.

Note: *The denominator of the C_{OUT} equation cannot be a negative value, which implies a negative capacitor (such as a battery).*

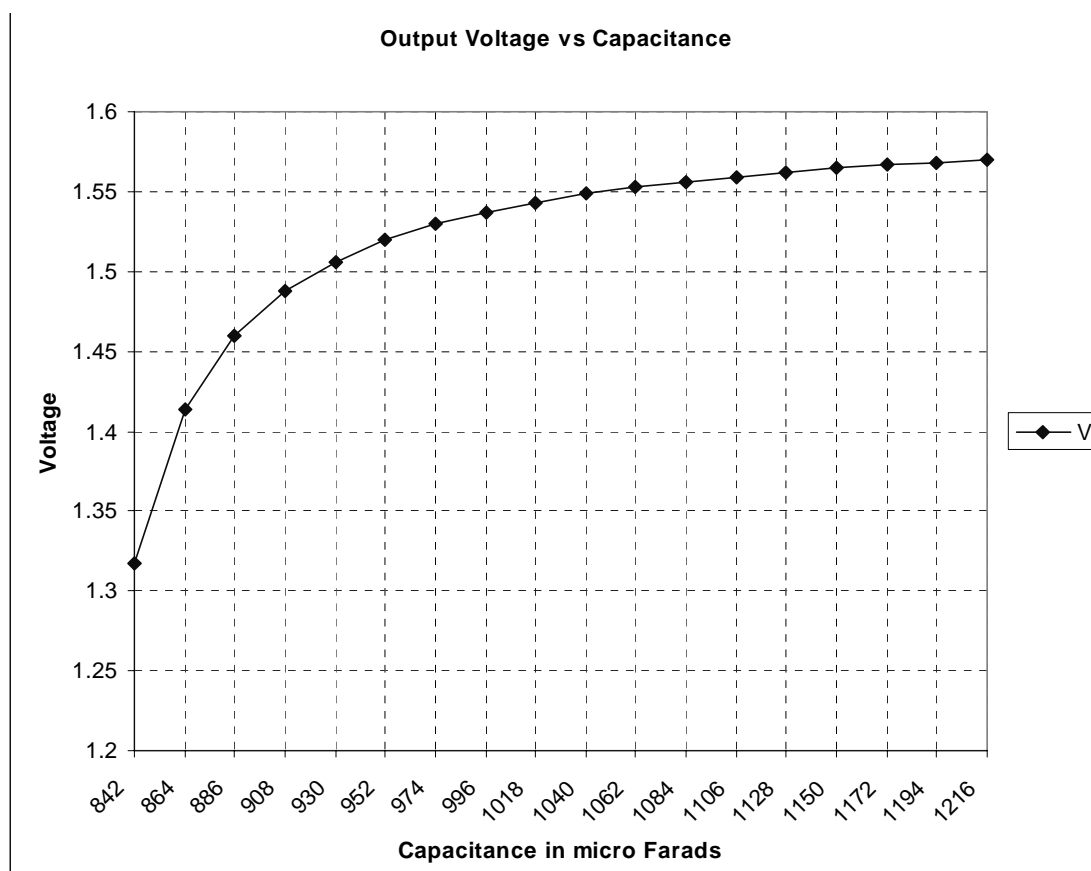


Figure 6. Bulk Decoupling versus Output Voltage Response for 1.6 V @ 35 A

Output Voltage Response Measurement Techniques

To measure output voltage response, run a program that creates a large current draw on the processor. AMD suggests using the popular benchmark program SPECfp under Microsoft® Windows NT® 4.0. While running such a program, toggle STPCLK# every 40 μ sec or less.

Measure the voltage differentially at the feedback pins on the slot connector, COREFB+ and COREFB-, located at pins A120 and A121, respectively. Use a scope probe with a ground connection next to the tip. The 3-inch to 6-inch ground leads that come off the side of a scope probe have too much inductance for this type of measurement. The scope bandwidth should be limited to 20MHz, giving a clear indication of the power supplied and bulk decoupling verification.

AMD used a Tektronix 684B scope with 6245 probes and an HP54720 with 54701 probes. (There was no significant difference between these two instruments.) The data was taken during a 40-second window with the scope set to infinite persistence. For a good starting point, use a horizontal sweep rate of 500nsec per division and a vertical scale of 0.1 V per division.

**Output Voltage
Response
Measurement
Utilities**

AMD has developed a maximum power utility to assist in designing systems that comply with the processor power and thermal requirements. This utility can verify that the supply voltage remains stable during a transition to a higher power and current consumption level.

This utility is DOS-based. For systems based on the Microsoft Windows® 98 operating system, re-boot in DOS mode or boot from a bootable DOS floppy disk that contains the utilities. For systems based on the Windows NT and OS/2 operating systems, boot from a bootable DOS floppy disk that contains the utilities.

The maximum power utility is available under a nondisclosure agreement. Contact the AMD sales office for information.