

How (Not) to Decouple High-Speed Operational Amplifiers

Bruce Carter

High Performance Linear Products

ABSTRACT

Decoupling the power supply pins of high-speed operational amplifier circuits is critical to their operation. Decoupling is also one of the least understood topics in engineering. It is seldom given the time or care required, yet it is relatively simple. This document will explain the pitfalls in decoupling and offer some correct techniques.

Introduction

Decoupling in high-speed design is usually done with little care. Engineers are so eager to get a prototype operational that they grab a handful of 0.1- μF or 0.01- μF capacitors out of a laboratory bin and assume, whew, the job is done. They think they are good engineers who have *solved* the problem of high-frequency coupled noise by throwing capacitors (and money) at it. They placed the obligatory capacitors, much like paying obligatory taxes to the government.

Decoupling is a design task worthy of at least the same degree of analysis as calculating an operational amplifier's gain or filter. Proper decoupling techniques do not have to be any more difficult than other design tasks. It requires the discipline on the part of the designer to take the task seriously and spend the time required to do it properly. A good designer should resist the temptation to grab capacitors at random from a laboratory bin or from a reel on the production floor.

When a capacitor is used for decoupling, it is connected as a shunt element to carry RF energy from a specific point in a circuit, away from a circuit power pin, and to ground. Ideally, the impedance of the capacitor to ground should appear as low as possible to the RF energy that needs to be rejected. It is important to know the frequencies that are causing problems and to select the right capacitor(s) to eliminate those frequencies.

Know Your Capacitor!

All capacitors are not created equal. A 0.1- μF or 0.01- μF capacitor from a lab bin may or may not be what is expected. Most lab stock is stripped of information that would indicate the part's manufacturer, grade, dielectric used, or its frequency characteristics. At high frequencies, the capacitor may operate quite differently from the *brick-wall high-frequency short* that the designer wants.

Capacitor Impedance and Series Self Resonance

Misconception number 1: Most designers think that a capacitor's value is independent of frequency. It is not. The published value of capacitance is taken at a frequency the manufacturer specifies. It is relatively constant at low frequencies, but the value can change by more than an order of magnitude at high frequencies.

Capacitors have a series self-resonant frequency due to parasitic inductance. Capacitors also have parallel self-resonances that will be discussed later. As the operating frequency approaches the capacitor's series self-resonant frequency, the capacitance value will appear to increase. This results in an effective capacitance C_E that is larger than the nominal capacitance.

The reason that capacitors are self-resonant is not hard to understand. A simplified high-frequency model of a capacitor includes inductive components:




Figure 1. High-Frequency Capacitor Model

The model terms in Figure 1 above are:

C_O : Nominal capacitance at a frequency defined by the manufacturer

L_S : Equivalent series inductance

ESR: Equivalent series resistance

R_P : Parallel (leakage) resistance

The effective capacitance C_E is a function of the reactance developed between the capacitor and its parasitic series inductance L_S . The equivalent series resistance ESR and leakage resistance R_P shown in Figure 1 do not have a significant effect on the effective capacitance. They do, however, affect the Q of the series self-resonance of the device. The effective capacitance C_E of a capacitor is defined as:

$$C_E = \frac{C_O}{1 - (2\pi F_O)^2 L_S C_O}$$

A capacitor's series self-resonant frequency F_{SR} is defined as:

$$F_{SR} = \frac{1}{2\pi\sqrt{L_S C_O}}$$

The Q, or sharpness of the resonance characteristic, is given by the expression:

$$Q = \frac{|X_C - X_L|}{ESR \parallel R_P}$$

At the capacitor's series self-resonant frequency, the reactance from C_0 and L_S are equal and opposite, yielding a net reactance of zero. At this self-resonance, the net impedance will be equal to $ESR \parallel R_p$. The capacitance at the series self-resonance will therefore be undefined, and the capacitor will have the lowest impedance at the series self-resonant frequency when used for decoupling. This impedance is typically in the range of milliohms ($m\Omega$). At frequencies below series self-resonance, the impedance will be capacitive, and at frequencies above series self-resonance, the impedance will be inductive. The impedance of the capacitor increases rapidly as the frequency is moved away from series self-resonance, and is given by the expression:

$$Z_C = \sqrt{(ESR \parallel R_p)^2 + (X_L - X_C)^2}$$

Figure 2 shows the response for two typical capacitors.

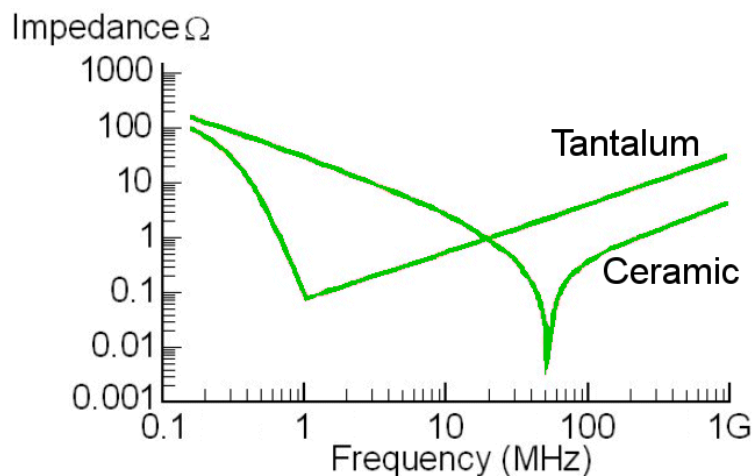


Figure 2. Typical Series Self-Resonances of Capacitors

The series self-resonant frequencies are relatively low. This should be a clue to high-frequency designers that they need to take great care in selecting bypass capacitors. Some hints follow:

- A designer can maximize rejection of an unwanted RF frequency by selecting a capacitor with a series self-resonant frequency corresponding to the unwanted frequency.
- No single capacitor can be expected to properly decouple a system where multiple RF frequencies are present. Capacitors can be placed in parallel to reject different RF frequencies.
- A capacitor becomes a dc blocking inductor above the series self-resonant frequency, which can present high-impedance to the RF interfering frequency, defeating the purpose of having a decoupling capacitor. This is why it is so bad to put the standard 0.1- μ F or 0.01- μ F capacitor (of questionable vintage) on a printed circuit board (PCB). It lulls the designer into a false sense of security, thinking that decoupling is done, when in reality all that capacitor is doing is taking up room on the PCB.

Most capacitor manufacturers give the series self-resonance characteristic of their capacitors as a graph of frequency vs capacitance. Multiple lines are sometimes placed on the chart, corresponding to different capacitor package sizes.




Figure 3. Bypass Capacitor Selection

It is important to be close: Q values for microwave capacitors can be in the hundreds, meaning that a small error in series self-resonant frequency can change the attenuation by orders of magnitude.

To select the correct capacitor to reject a given frequency, follow the frequency line on the chart over to where it intersects the line on the graph, then go down the chart to the capacitance value. In Figure 3, to reject 1.9 GHz, 15 pF would be the correct value of the capacitance.

This is a far cry from mindlessly selecting a 0.01- μ F or 1000 pF capacitor, and it illustrates another misconception.

Misconception number 2: Most designers think that when poor decoupling is suspected, increasing the value of capacitance will produce more attenuation. The exact opposite is true. When poor decoupling of high frequencies is suspected, decrease the value of capacitance (or preferably, do the homework and look up the series self-resonance).

Correctly bypassing a high-speed operational amplifier circuit is as simple as maintaining a prototyping kit of microwave capacitors and looking up the correct value(s) of capacitance on a graph.

If more than one frequency (or band of frequencies) is present in the system, two or more capacitors may be required to properly bypass. Each capacitor is then targeted at a different frequency range.

Parallel Self-Resonance

Capacitors also have parallel self-resonance frequencies, related to the way they are manufactured and mounted on a PCB. There is only one series resonant frequency, but there are a series of parallel capacitances. A rough rule of thumb is that the first parallel resonant frequency occurs at about twice the series resonant frequency.

While a capacitor has its lowest impedance at the series self-resonant frequency, it typically has its highest impedance at the parallel self-resonant frequencies. Therefore, while it is possible to take advantage of the series resonant frequency for RF purposes, the capacitor is useless at parallel resonant frequencies.

It is not hard to understand the mechanism behind parallel resonance. Figure 4 shows a cutaway picture of a capacitor on a PCB.




Figure 4. Capacitor on a PCB

The capacitor is a collection of parallel plates. These plates sum together to produce the effective capacitance C_E , but because all conduction is skin conduction at high frequencies, there is also some parasitic parallel capacitance plate-to-plate. In addition, because most capacitors are mounted horizontally on PCBs, there is also parasitic capacitance from the horizontal plates to the ground and power planes of the PCB.

The high-frequency equivalent model is modified to include parallel resonant effects:




Figure 5. High-Frequency Capacitor Model with Parallel Capacitance

C_P in this schematic is actually a lumped element corresponding to the individual contributions from each plate and the PCB. The overall response of this circuit is shown in Figure 6. The resonances are shown as nulls, the magnitude of which is affected by source and load impedance. Vertical scale has been purposely omitted, as it will vary greatly in each application depending on the load. The presence of a null locates a resonance frequency.




Figure 6. Response of a High-Frequency Capacitor Above Series Self-Resonance

The first parallel resonant frequency (PRF1) occurs very close to the series self-resonance (SRF). This is a danger sign to the designer that capacitor tolerance could change the attenuation at SRF into transmission at PRF1 very easily. It is very desirable to get rid of PRF1!

Fortunately, there is a very easy way to eliminate PRF1, as well as all odd parallel resonant frequencies. If the capacitor is mounted on its side, the plates inside the capacitor are now perpendicular to the PCB:




Figure 7. Vertical Capacitor Mounting

This will eliminate the parasitic capacitances between the plates of the capacitor and planes on the PCB. The resulting spectrum shows that PRF1, as well as all odd parallel resonances, are eliminated:




Figure 8. Resulting Response from Mounting a Capacitor Vertically

The response curve above shows that the capacitor will be usable for a much larger range of frequencies above the SRF, providing that predominately-inductive impedance is acceptable. The first region of high impedance at PRF2 does not occur until almost 2.5 times the SRF frequency.

Some microwave capacitor manufacturers, such as Johanson, supply software to plot the self-resonances of their products. These can greatly aid in designing a high-speed decoupling application, or any high-speed or RF application where the capacitor characteristics will be important.

Mounting a capacitor in this way runs counter to the training of manufacturing and inspection personnel, who will have to revise their thinking for a high-speed PCB.

Capacitor Dielectric Types

Just as the internal construction and orientation of capacitor plates affect its self-resonant characteristics, the material between the plates, the dielectric, affects the characteristics of the capacitor. Dielectric material is the medium that isolates the plates, and it has characteristics of its own that affect the capacitor performance. The following information is condensed from one manufacturer:

Table 1. Capacitor Dielectric Characteristics

Dielectric Name	Stability	EIA Class	TC	Range
NPO or COG	Very stable, high Q	I	± 30 PPM	-55 to +125 °C
X7R	Semistable	II	±15%	-55 to +125 °C
Z5U		III	+22 / -56%	+10 to +85 °C
Y5V		III	+22 / -82%	-30 to +85 °C

The average bin of laboratory stock capacitors is probably X7R or worse, because these are lower cost than the very desirable NPO (also called COG). 0.1- μ F or 0.01- μ F value capacitors used as decoupling capacitors are relics of the age of discrete digital design. The author has examples of early digital design from the 1960s with 0.1- μ F ceramic disk capacitors lined up in neat rows between rows of digital ICs. Printed on every capacitor is the tolerance +20 / -80%, showing that the dielectric type is Y5V. The clock speed of these early digital boards, however, was a few megahertz at most. Nevertheless, 0.1- μ F Y5V capacitors have persisted to the present day as decoupling capacitors, even as operating speeds for both digital and analog systems have increased a thousand fold. It is time for designers to adopt a new mindset about decoupling. The smartest decoupling design decision a designer can make would be to take the existing lab stock of supposed decoupling capacitors and dump them into a waste container. High-frequency decoupling requires NPO/COG capacitors.

NPO/COG capacitors tend to get expensive at high values of capacitance. Fortunately for designers, it will seldom be necessary to use high values of capacitance.

What About Electrolytic Capacitors?

Notice that nowhere in the preceding discussion was anything said about electrolytic capacitors. There is a good reason for this. Electrolytic capacitors, whether they are aluminum or tantalum, are very-low-speed devices. Their self-resonant frequency is limited to a range between 100 kHz and 1 MHz. Therefore, they are no good for decoupling high frequencies. The primary purpose of electrolytic decoupling capacitors is to filter power-supply switching noise, which is in this range.

Electrolytic capacitors are often seen in digital systems, which are often powered from switching power supplies. Large digital systems often require many amperes of current, and the power supply has severe constraints. Switching power supplies are small, light weight, efficient, and therefore produce a lot of wattage in a small volume without generating a lot of heat. The only disadvantage of switching power supplies is ripple and noise. Ripple, of course, is the artifact of 50 Hz or 60 Hz. Noise is an artifact of ac-to-dc conversion in the front end of the supply, which tends to be quite noisy.

The primary use, then, for electrolytic capacitors, is at frequencies below 1 MHz. If a designer is sure that these frequencies are present, then they are necessary. If, however, there is no reason to believe that these frequencies are present, electrolytic capacitors are not necessary. A good example would be in a cellular telephone. The telephone is powered off a battery when operated, and it is operated off a linear power supply when recharging. If no dc-to-dc conversion is done inside the unit, then no electrolytic capacitor would be needed for suppressing switching power supply noise. If a dc-to-dc converter is used, it may be operating at frequencies higher than 1 MHz, and an electrolytic will not work. If, however, the telephone is operated near an AM radio transmitter, an electrolytic capacitor may be necessary to suppress RF that comes in through a *hands-free* device. This type of system level analysis and scenario development is needed for decoupling decisions.

Texas Instruments includes electrolytic capacitors on evaluation boards because there is a chance they will be operated off a switching power supply, or operated near AM stations with long power-supply leads. Most laboratory applications, however, are well-shielded from AM radio stations, and most laboratory power supplies are not switching supplies. Therefore, electrolytic capacitors are not usually required.

It is advantageous to omit electrolytic capacitors whenever possible. There are a couple of reasons:

- Most electrolytic capacitors are on allocation from manufacturers, with long lead times.
- Electrolytic capacitors are large, and they force other traces in a high-speed design to be correspondingly longer to accommodate them. High-speed design requires that traces be as short as possible to avoid parasitic capacitance and inductance. A component that is not absolutely required should not be on the board displacing traces, particularly a large component such as an electrolytic capacitor.
- The pads required to support an electrolytic capacitor disrupt the width and placement of traces and vias, creating an opportunity for RF noise to be radiated into the circuit at the point where the capacitor is placed.

How to Lay Out Decoupling Capacitors on the PC Board

Once the proper capacitor has been selected, it must be applied properly on the PCB, or all of the care in selection will be in vain.

Trace Inductance

Semiconductor manufacturers are responsible, in part, for poor high-frequency design techniques. Consider the DIP 14 IC package. The power pin has been placed on the diagonal opposite corner from the ground pin. There are a limited number of ways to make the connection to a decoupling capacitor:




Figure 9. Traditional Decoupling Layouts

- If the capacitor is placed near the power pin, there is a long connection to the ground pin, and access to the left side of the IC is restricted on one layer.
- If the capacitor is placed near the ground pin, there is a long connection to the power pin, and access to the right side of the IC is restricted on one layer.
- If the connections are made underneath the IC, there is free access to both sides of the IC on all layers, but access beneath the IC is blocked, and nothing is done to alleviate long connections.
- The capacitor can be placed on the bottom side of the PCB (not shown), which will produce the shortest possible traces, providing that orthogonal routing is abandoned. The traces, however, would still add to more than a centimeter.
- The worst problem, by far, created by the long connections is trace self-inductance. The formula for the inductance of a wire or PCB trace can be approximated by:

$$L(\text{nH}) = 2x \cdot \left(\ln \left(\frac{2x}{w+h} \right) + 0.2235 \left(\frac{w+h}{x} \right) + 0.5 \right)$$

Where:

x = length of the trace (cm)

w = width of the trace (cm)

h = height of the trace (cm)

The inductance is relatively unaffected by the height of the trace. It is more affected by the width, but it takes a large change in width to substantially affect the inductance. The predominant effect is the length. Common PCB traces have self-inductances that measure between 6 nH and 12 nH per centimeter.

This is very bad news at high frequencies. For example, assume:

- A cell phone application that must decouple 1.9 GHz
- 6 nH/cm trace self inductance
- 1.5 cm of circuit traces to the decoupling capacitor
- A 30-pF capacitor with a C_E of 27.77 pF, and an L_S of 0.23 nH

The first step is to use the series self-resonant frequency model and calculate the SRF of the capacitor, which will be 1.92 GHz, ideal for cell phones. But wait! Inductance adds in series. If 9 nH is added to the L_S of the capacitor, the series self-resonance becomes 314 MHz. This PCB layout is entirely unsuitable. Clearly, trace inductance must be avoided at all costs! It changes the series self-resonance of the capacitor.

The example above is a bit extreme. Very few designers would think to use a DIP-14 IC in a cell phone. Newer ICs are surface-mount, and many have moved the power pins to more convenient locations. The movement to surface-mounting helps somewhat, because the connection geometries shrink. But shrinking geometries alone do not help the problem. The self-inductance of PCB traces is so large that any trace length negates careful capacitor selection. Fortunately, most high-frequency PCBs are now multilayer and have both power and ground planes. Designers should place vias from the planes as close as possible to the pad of the capacitor. Preferably, they should be incorporated into the copper structure of the pad itself.




Figure 10. Proper Feedthru Method

Even a couple of millimeters, as shown in the *bad* example in Figure 10, adds 1.2 nH to the L_S of the capacitor. Using PCB traces to make connections to decoupling capacitors simply cannot work! The best that the designer can achieve is to minimize trace length by moving the feedthrus as near as possible to the part and making the resulting connection as wide as possible.

Via Inductance

Whenever there is a via on a PCB, a parasitic inductor is also formed. At a given diameter (d) the approximate inductance (L) of a via at a height of (h) may be calculated as follows:

$$L \approx \frac{h}{5} \cdot \left(1 + \ln \left(\frac{4h}{d} \right) \right) \text{ nH.}$$




Figure 11. Cutaway View of Inductive Vias

For example:

A 0.4-mm diameter via through a 1.5-mm thick PCB has an inductance of 1.1 nH. This is also extremely bad news for high-frequency designers.

Three techniques combat via inductance:

- Many PCB manufacturers offer laser drill systems that can create microvias with diameters from 0.15 mm down to 0.025 mm.
- A blind via will reduce via height and, therefore, the inductance.
- Placing multiple inductors in parallel can reduce inductance.

If the example above is changed to:

- Two blind microvias with a diameter of 0.025 mm going to an adjacent layer on a ten-layer board ($h = 0.15$ mm), the via inductance is reduced to 0.063 nH.

Height is the predominant factor, so blind vias are much more important than microvias. If the example is changed to:

- One blind microvia with a diameter of 0.15 mm going to an adjacent layer on a ten-layer board ($h = 0.15$ mm), the via inductance is reduced to 0.072 nH.

This may reduce PCB cost significantly, because 0.15 mm is the lower limit of conventional drilling techniques, making the laser-processing step unnecessary. It may also save the cost of special plating that is capable of plating a 0.025 mm hole.

An Alternative Approach

If blind vias and laser processing are not acceptable for cost reasons, the designer sometimes can borrow layout techniques from RF designers to combat parasitic inductances. These techniques do not necessarily suppress parasitic inductance, but they place the parasitic inductances where they do not affect circuit performance.




Figure 12. High-Speed Layout with RF Techniques

Figure 12 has several features that might escape the casual observer.

- It is a multilayer board, with internal power and ground planes.
- Multiple connections are made with vias around the outside of the circuit to a ground plane. The parasitic inductance is low because it adds in parallel, reducing the overall value.
- Signal pins are connected to vias, to components on the backside of the board. In some applications, small circuits will probably be constructed on the top layer.
- Ground-plane connections to the IC and the decoupling capacitor are made on the top layer with thermal reliefs. Thermal reliefs do act as traces and, therefore, add a small amount of inductance, but they are necessary for proper soldering.
- Connection to power is made with a via, but the via and trace inductances appear before the capacitor, which is on the same pad as the power pin of the IC. Therefore, the parasitic via and trace inductance do not add to the capacitor's series self-resonance. They actually help, by adding impedance at high frequencies to the power connection.
- Remember that the capacitor should be mounted vertically on its side.
- The capacitor is also connected to the power pad of the IC through thermal reliefs.
- When power and ground planes are used in a multilayer board, they add distributed capacitance to the system, improving broadband noise rejection.

Conclusions

Decoupling high-speed operational amplifier circuitry requires a little bit of design work. Capacitors should be selected for their high-speed characteristics, including self-resonant frequency. They should be mounted vertically (on the side) to avoid parasitic capacitance with the PCB. Care should be taken when using PCB traces and vias, because both have parasitic inductance that adds to that of the capacitor.

Bibliography

1. *Capacitors in Bypass Applications*, Richard Fiore, American Technical Ceramics, www.atceramics.com
2. *Effective Capacitance vs Frequency*, Richard Fiore, American Technical Ceramics, www.atceramics.com
3. *Decoupling Basics*, Arch Martin, AVX Corporation
4. *SRF & PRF and Their Relation to RF Capacitor Applications*, Johansen Technology
5. *Ask The Applications Engineer*, James Bryant, Analog Devices

IMPORTANT NOTICE

Texas Instruments Incorporated and its subsidiaries (TI) reserve the right to make corrections, modifications, enhancements, improvements, and other changes to its products and services at any time and to discontinue any product or service without notice. Customers should obtain the latest relevant information before placing orders and should verify that such information is current and complete. All products are sold subject to TI's terms and conditions of sale supplied at the time of order acknowledgment.

TI warrants performance of its hardware products to the specifications applicable at the time of sale in accordance with TI's standard warranty. Testing and other quality control techniques are used to the extent TI deems necessary to support this warranty. Except where mandated by government requirements, testing of all parameters of each product is not necessarily performed.

TI assumes no liability for applications assistance or customer product design. Customers are responsible for their products and applications using TI components. To minimize the risks associated with customer products and applications, customers should provide adequate design and operating safeguards.

TI does not warrant or represent that any license, either express or implied, is granted under any TI patent right, copyright, mask work right, or other TI intellectual property right relating to any combination, machine, or process in which TI products or services are used. Information published by TI regarding third-party products or services does not constitute a license from TI to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property of the third party, or a license from TI under the patents or other intellectual property of TI.

Reproduction of TI information in TI data books or data sheets is permissible only if reproduction is without alteration and is accompanied by all associated warranties, conditions, limitations, and notices. Reproduction of this information with alteration is an unfair and deceptive business practice. TI is not responsible or liable for such altered documentation. Information of third parties may be subject to additional restrictions.

Resale of TI products or services with statements different from or beyond the parameters stated by TI for that product or service voids all express and any implied warranties for the associated TI product or service and is an unfair and deceptive business practice. TI is not responsible or liable for any such statements.

TI products are not authorized for use in safety-critical applications (such as life support) where a failure of the TI product would reasonably be expected to cause severe personal injury or death, unless officers of the parties have executed an agreement specifically governing such use. Buyers represent that they have all necessary expertise in the safety and regulatory ramifications of their applications, and acknowledge and agree that they are solely responsible for all legal, regulatory and safety-related requirements concerning their products and any use of TI products in such safety-critical applications, notwithstanding any applications-related information or support that may be provided by TI. Further, Buyers must fully indemnify TI and its representatives against any damages arising out of the use of TI products in such safety-critical applications.

TI products are neither designed nor intended for use in military/aerospace applications or environments unless the TI products are specifically designated by TI as military-grade or "enhanced plastic." Only products designated by TI as military-grade meet military specifications. Buyers acknowledge and agree that any such use of TI products which TI has not designated as military-grade is solely at the Buyer's risk, and that they are solely responsible for compliance with all legal and regulatory requirements in connection with such use.

TI products are neither designed nor intended for use in automotive applications or environments unless the specific TI products are designated by TI as compliant with ISO/TS 16949 requirements. Buyers acknowledge and agree that, if they use any non-designated products in automotive applications, TI will not be responsible for any failure to meet such requirements.

Following are URLs where you can obtain information on other Texas Instruments products and application solutions:

Products		Applications	
Amplifiers	amplifier.ti.com	Audio	www.ti.com/audio
Data Converters	dataconverter.ti.com	Automotive	www.ti.com/automotive
DSP	dsp.ti.com	Broadband	www.ti.com/broadband
Interface	interface.ti.com	Digital Control	www.ti.com/digitalcontrol
Logic	logic.ti.com	Military	www.ti.com/military
Power Mgmt	power.ti.com	Optical Networking	www.ti.com/opticalnetwork
Microcontrollers	microcontroller.ti.com	Security	www.ti.com/security
RFID	www.ti-rfid.com	Telephony	www.ti.com/telephony
Low Power Wireless	www.ti.com/lpw	Video & Imaging	www.ti.com/video
		Wireless	www.ti.com/wireless

Mailing Address: Texas Instruments, Post Office Box 655303, Dallas, Texas 75265
Copyright © 2007, Texas Instruments Incorporated