The MSP430x3xx Clock System

Darren Staples *Mixed Signal Products*

ABSTRACT

This application report describes in detail the oscillator system present on the MSP430x3xx devices, including use of the digitally-controlled oscillator (DCO), the frequency locked loop (FLL), and the low-power modes.

Contents

List of Figures

List of Tables

1 Introduction

The MSP430 family of microcontrollers offers designers a unique blend of 16-bit processing power, flexible and easy to use peripherals, and ultra-low power consumption. In a typical 3-V application running at 1 MHz, single-cycle 16-bit instructions are executed at a speed of one million instructions per second, with a current consumption of only 400 μA. True, real-time interrupt handling can now be achieved with the MSP430's rapid wake up from low power mode 3 (only 1.5 μA of current consumption) to fully-active mode in only 6 μs.

These remarkable capabilities are the direct result of two of the MSP430's most interesting design features—its 16-bit RISC-like architecture and its clock system.

NOTE: The MSP430x3xx clock system is significantly different from the clock system present in the MSP430x1xx devices. This document will cover only the MSP430x3xx. Other application-note material exists for the MSP430x1xx.

2 The MSP430x3xx Family's Clock System

2.1 The 32-kHz Crystal Oscillator

The design of a microcontroller clock system is highly critical to the overall performance of an application. In order to be cost-effective, stable, and accurate, most designs use a quartz crystal or ceramic resonator as the frequency reference. These devices typically operate in the range of 100 kHz to 10 MHz. However, both of these choices exhibit disadvantages. The current consumption of the oscillator will rise in proportion to any rise in oscillation frequency, so higher frequencies are not advantageous for low-power designs. Also, when powered up, quartz-crystal oscillators take a significant period of time (in the region of tens to hundreds of milliseconds) to reach stable-oscillation frequency and amplitude due to the high-Q factors of the crystals utilized. This does not provide fast response to interrupt stimuli.

The MSP430x3xx avoids these problems by using an ultralow-power oscillator based around a 32-kHz quartz crystal of the kind commonly used in wrist watches and clocks. These crystals have a small physical size, and are readily available and cost-effective. Most importantly, the oscillation frequency is low enough to minimize current consumption. Low power mode 3 is a shutdown mode in which only the 32-kHz crystal oscillator is left active—a typical current consumption figure for this mode is 1.5 μA. The crystal oscillator can therefore be left running continuously, thus avoiding the start-up and stabilization delays mentioned above. Having a 32 kHz clock continuously available also means that some of the on-chip peripherals can be left active even when the rest of the system is shut down. For example, the liquid-crystal display can be left on, or one of the timers can be used to implement a real-time clock.

The MSP430x3xx's internal oscillator design means that no additional external capacitors are required for the 32-kHz crystal. However, the circuit traces connecting the crystal to the MSP430x3xx X_{IN} and X_{OUT} pins should be made as short as possible to minimize the effects of external interference.

The 32-kHz clock signal generated by the crystal oscillator is named ACLK, or *auxiliary clock*. This signal is distributed throughout the system for use by the on-chip peripherals.

2.2 The Frequency Multiplier

The ACLK alone can not enable the MSP430x3xx to achieve the processing throughput required, so a frequency multiplication technique is used for this purpose.

Traditionally, frequency multipliers have been based upon phase-locked loops, which can suffer from stability and time-to-lock problems when used over a wide frequency range. Also, the loop-control mechanism is generally an analog component requiring a constant power supply.

Instead, the MSP430x3xx frequency multiplier is based around a frequency-locked loop. This is a purely digital system, and consequently it can be controlled by software. Its current consumption is zero when inactive.

Figure 1 shows a block diagram of the frequency multiplier. The digitally-controlled oscillator, or DCO, generates an output clock frequency named MCLK or *master clock*. This frequency is divided by the software-controlled factor N. The synchronizer generates an *error signal* to the frequency integrator, based upon the difference between ACLK and the divided MCLK. The frequency integrator is basically an up/down counter which adjusts the DCO output frequency MCLK. Once frequency lock has been achieved, the synchronizer *error signal* becomes exactly zero, and MCLK becomes exactly N times ACLK.

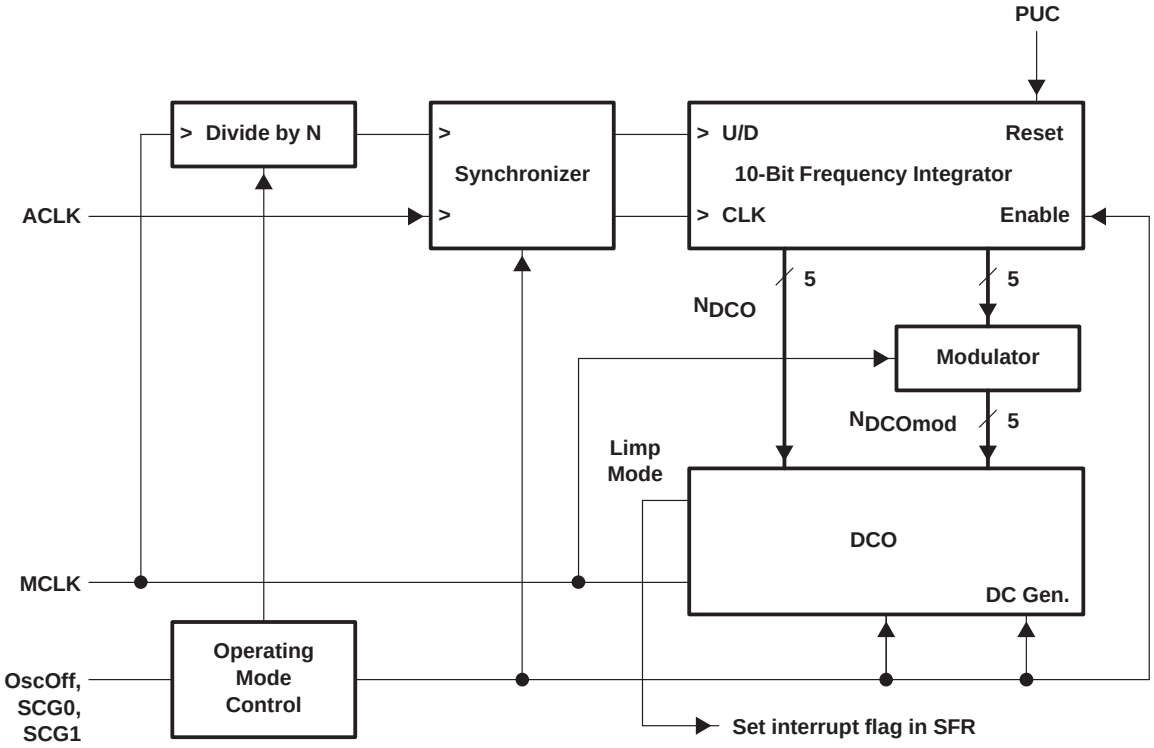

Figure 1. Frequency Multiplier Block Diagram

However, by its very nature, a DCO can generate only a limited number of discrete frequencies. There are only 32 frequencies or *taps* available on the MSP430x3xx, each being approximately 10% higher than the previous one. So, depending on the divide ratio N and various other factors, it is likely that the synchronizer error signal will never reach exactly zero, and MCLK will never be exactly N times ACLK.

2.3 The Modulator

In cases where the MCLK is not an exact multiple of ACLK, the modulator adjusts the control signals from the frequency integrator, and causes the DCO to hop between two adjacent frequency taps, one on either side of the required *fractional tap*, according to a predetermined pattern. Over a number of MCLK cycles it can be seen that the *average* synchronizer *error signal* will then be exactly zero, and the average MCLK frequency will be exactly N times ACLK.

Figure 2 shows a typical *fractional tap* between two discrete DCO output-frequency taps.

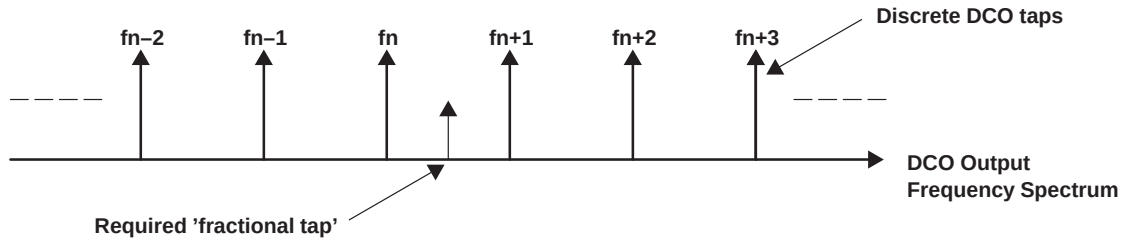

Figure 2. Fractional Tap Frequency Required

Figure 3 shows some of the 32 modulator *hop* patterns.

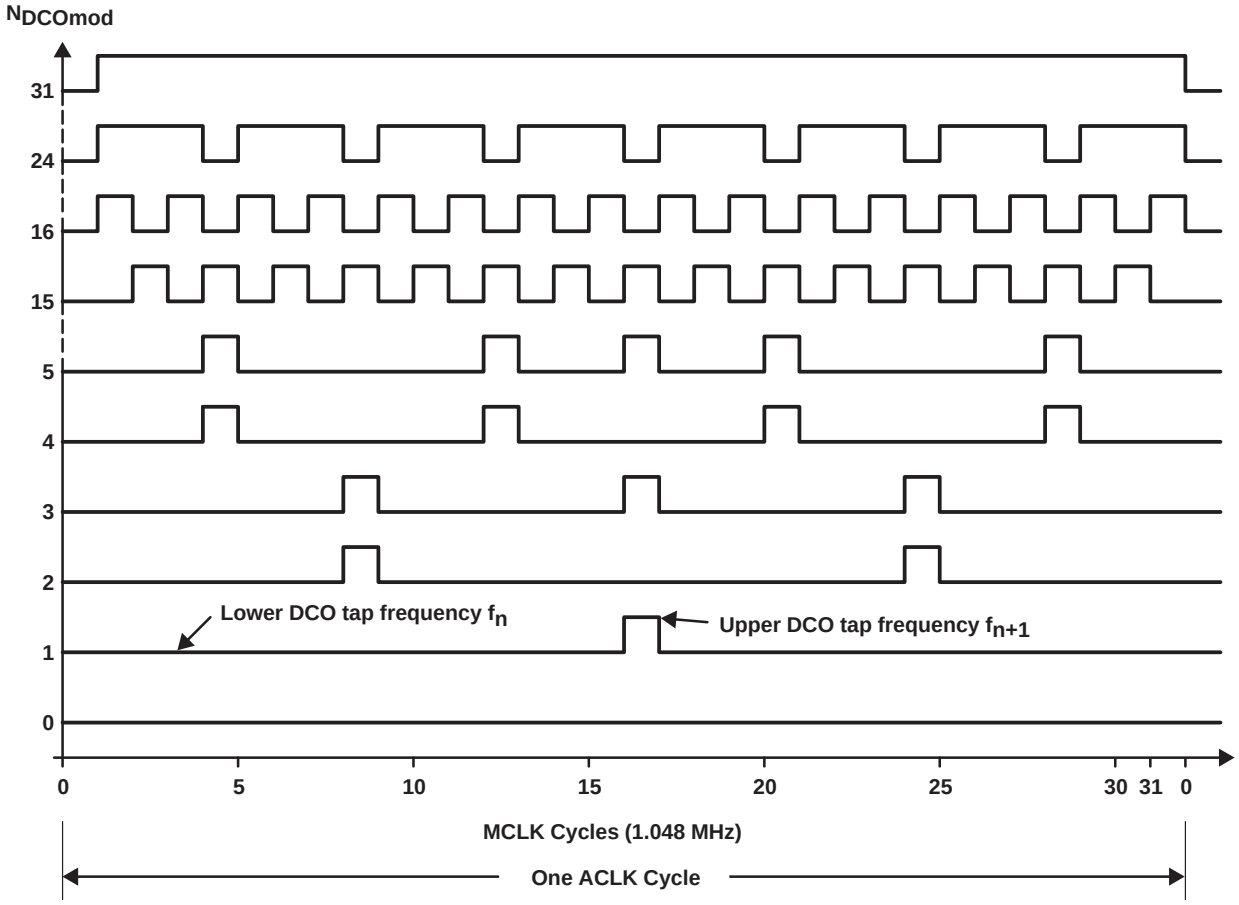

Figure 3. Modulator Hop Patterns

The frequency-locked-loop and modulator are enabled by default after a system reset. In the majority of simple applications, the system can be left in this state. The frequency-locked-loop compensates for temperature and operating voltage drifts that occur over time, and MCLK is accurately maintained.

This may not be ideal in all circumstances, however. Consider the case where an application needs to perform high-resolution timings. Serious errors can be introduced if the modulator adjusts the MCLK frequency midway through a critical timing sequence. It is therefore possible to disable both the frequency-locked-loop and the modulator.

With modulation enabled, MCLK is adjusted by the modulator, when necessary, potentially every cycle (see hop pattern number 16 in Figure 3). With modulation disabled, MCLK is adjusted only every 1,024 cycles. If a completely-stable MCLK is required, then both the modulator and the frequency-locked-loop must be disabled. The DCO then remains at its current frequency tap. However, for the MCLK to remain stable, the DCO should not be subjected to temperature and operating voltage drifts. In most practical situations, the period during which such a stable MCLK is required is relatively short compared to the period during which a significant temperature or operating voltage drift would occur. Once the MCLK stability is no longer required, it is recommended to reenable the modulation and the frequency-locked-loop to allow compensation for longer-term drifts. In fact, for any processing period longer than 100 MCLK-cycles, it is recommended to reenable the frequency-locked-loop and modulator.

It was previously stated that MCLK = N x ACLK, where N is a software-programmable factor. This multiplication factor N is controlled by the lower seven bits of the SCFQCTL register. Note that the value written to the SCFQCTL register is automatically incremented by one. In other words, to achieve a frequency-multiplication factor of 32, 31 should be written into the lower seven bits of SCFQCTL.

This ability to dynamically adjust the main processor clock frequency offers some interesting possibilities. For example, in a battery-powered system it means that the overall system performance can be continuously adjusted to suit current circumstances to get the absolute maximum lifetime from the batteries.

Consider the case where an application is able to perform an extended period of processing, during which current consumption must be minimized. By reducing the factor N, MCLK can be reduced. The processor core's current consumption is directly proportional to MCLK; consequently, this current is also reduced. Alternatively, if a short period of intense processor activity is required, where current consumption is not critical, MCLK can be increased. Once a normal level of processing is required, MCLK can be returned to its nominal level.

The five N_{DCO} bits select which of the 32 DCO taps will be used. These bits are contained in register SCFI1. The five N_{DCOMOD} bits control the modulator and are shared between the SCFI0 and SCFI1 registers. Both N_{DCO} and N_{DCOMOD} can be read by the application software to see which frequency tap the DCO is currently using. Interestingly, these bits can also be written.

It should be noted that the DCO can generate maximum frequencies beyond the specified operating ranges of the CPU and peripherals; so care should be taken when programming the frequency multiplication factor N.

At reset, the default value of N is 32, to give a target MCLK frequency of 1.048 MHz. Immediately after reset, however, the DCO's lowest tap frequency is automatically selected, giving an MCLK frequency of approximately 500 kHz. Code execution begins immediately after the reset, and all peripherals are then available for use at this reduced MCLK frequency. The frequency-locked-loop then adjusts MCLK towards the target frequency by stepping through the frequency taps. To step from one tap to the next requires 1,024 MCLK cycles, and in the worst case, the DCO may have to step through all 32 taps to reach the target frequency. Therefore, the application software should not perform any timing operations which rely on an accurate MCLK until at least 32,768 MCLK cycles have elapsed since the reset.

2.4 DCO Range Control

When MCLK is changed by adjusting the frequency multiplication factor N, the frequency-locked-loop adjusts the DCO frequency towards the desired target frequency, and each step to the next DCO tap incurs a 1,024-cycle delay before MCLK reaches stability at the new tap frequency. It can be seen that, for large changes in MCLK frequency, significant stabilization delays could be introduced. So the MSP430x3xx has an alternative mechanism for making large MCLK frequency changes. The clock-control register SCFI0 contains three bits, FN_2, FN_3, and FN_4, which can be thought of as range-control bits for the DCO. By setting these bits to the appropriate pattern, as shown in Table 1, the nominal MCLK frequency (1.048 MHz) can be doubled, tripled, or quadrupled.

Table 1. The DCO Range Control Bits

FN_4	FN_3	FN_2	MCLK FREQUENCY
0	0	0	1 x $f_{nominal}$
0	0	1	2 x $f_{nominal}$
0	1	X	3 x $f_{nominal}$
1	X	X	4 x $f_{nominal}$

These range-control bits affect the DCO output frequency (MCLK) without changing the current DCO frequency tap (in reality, one of the immediately-adjacent taps may be selected). Hence, the time required for the DCO to settle at the new frequency is much shorter than the time when only the frequency-multiplication factor is adjusted.

Example: Doubling the MCLK frequency to 2.096 MHz

```
MOV.B #63, &SCFQCTL     ;MCLK = (63+1) * ACLK = 2.097 MHz
MOV.B #FN_2, &SCFI0     ;Set FN_2, for 2x DCO frequency
```

2.5 Low-Power Modes 3 and 4

When the MSP430x3xx enters low-power mode 3 (LPM3), the frequency integrator is not reset; that is, the DCO remains set at its current tap frequency, and the 32-kHz oscillator is left running. Consequently, when an interrupt occurs and the system reverts back to active mode, the DCO can start immediately and full speed operation begins within only 6 μs. However, the frequency-locked-loop is disabled by default. If the period of CPU activity is expected to exceed 100 clock cycles, then it is advisable to reenable the DCO loop control.

Current consumption in LPM3 mode is typically 1.5 μA, and the 32-kHz oscillator activity means that peripherals can be left running even when the rest of the system is shut down if they use the 32 kHz as their clock source. For example, a real-time clock could be implemented using one of the peripheral timers, or the LCD could be displaying user information even though the CPU and other peripheral devices are inactive.

Low-power mode 4 (LPM4) offers the lowest-current consumption of all the low-power modes—typically 0.1 μA. However, LPM4 is intended to be used primarily as a *storage and transportation mode* and must be used with some care due to the fact that RAM contents are retained as long as the supply voltage remains above the specified minimum. To achieve such a low-current consumption level, all the on-chip devices, including the 32-kHz oscillator, are shut down. This means that none of the on-chip peripherals can be left active as they can in LPM3. It also means that the system takes much longer to return to active mode than it does in LPM3 mode, because the 32-kHz oscillator has to be restarted as it happens after a reset. The quartz crystal itself can take hundreds of milliseconds to achieve stable oscillation frequency and amplitude.

There is an important difference in the-start up behavior of the 32-kHz oscillator between a reset and a return from LPM4. The frequency integrator is not reset when returning from LPM4, so the DCO remains set at the same frequency it was when the system went into LPM4. This can cause problems under certain circumstances. Consider the case where an application is about to go into LPM4. If the ambient temperature is high, then the DCO will be set to one of its upper-frequency taps, since it has a negative temperature coefficient. After a period of inactivity in LPM4, the system is reactivated when ambient temperature is much lower. The DCO will remain set at the upper frequency tap, and MCLK could be too high for the CPU or peripherals to cope with. It is therefore recommended that the lowest DCO tap be selected as part of the procedure for shutting down into LPM4.

LPM3 offers many advantages over LPM4, with the only disadvantage of a marginal increase in current consumption.

2.6 Oscillator Fault Detection

The MSP43x3xx clock system has a mechanism for warning the application when the DCO is operating outside its normal range. If the DCO is operating at one of the taps 0, 28, 29, 30, or 31, then an oscillator-fault interrupt can be generated. The oscillator-fault interrupt shares an interrupt vector with the non-maskable-interrupt (when the RST/NMI pin is configured as NMI). The oscillator-fault interrupt has an enable bit (OFIE) and an interrupt-flag bit (OFIFG). At reset, the OFIE bit is reset, disabling the oscillator-fault interrupt; but the OFIFG remains unchanged. Because the DCO starts from the lowest tap (tap 0) after a reset, the oscillator-fault flag OFIFG is automatically set. It remains set even when the DCO is operating at the central frequency taps. If the application software is designed to use the oscillator-fault functionality, then it must reset the OFIFG flag when the DCO is known to be operating at one of the *safe* tap frequencies (taps 1 to 27). Then It can also enable the interrupt by setting the OFIE bit. The oscillator-fault interrupt is non-maskable in the sense that the global interrupt enable bit (GIE, in the status register) does not affect it, even though the OFIE bit does. When an oscillator-fault interrupt occurs, the OFIE bit is automatically reset to prevent repeated interrupts. The application's oscillator-fault handler should then reenable the interrupt as required.

2.7 Buffered Clock Output

The MSP430x3xx provides a buffered oscillator output on the XBUF pin. This is intended for use in applications where external circuitry requires a clock. By programming the CBCTL register, the XBUF output can be configured to output ACLK, ACLK/2, ACLK/4, or MCLK, or it can be disabled.

2.8 Using the MSP430x3xx Without a 32-kHz Crystal

So far, this document has considered the MSP430x3xx clock system used with a 32-kHz crystal. The need for the crystal exists only if the application requires a frequency reference for accurate timing. In very simple applications without such timing requirements, it is quite feasible to let the DCO run freely.

Without a crystal, the DCO runs at its lowest-frequency tap by default: approximately 500 kHz. If this MCLK frequency is too low to be useful, then a higher frequency tap can be selected by programming the SCFI0 and SCFI1 registers. When doing this it should be noted, however, that the frequency-locked-loop should be disabled; otherwise the DCO is automatically stepped back down to 500 kHz.

Some applications can have other frequency references available. For example, an electricity meter can detect the zero-crossings of the power line's alternating voltage and, by using a simple timing mechanism, periodically adjust the DCO frequency so that a defined number of MCLK cycles occur for every voltage cycle. Other application-note material is available offering examples of this technique.

2.9 Clock System Control Registers

2.9.1 Status Register, SR, R1

Bit 7. SCG1.
Bit 6. SCG0.
Bit 5. OscOff.
Bit 4. CpuOff.

SCG1, SCG0, OscOff, and CpuOff control the operating mode of the MSP430x3xx devices according to the truth table shown in Table 2.

Table 2. The Operating Modes

SCG1	SCG0	OscOff	CpuOff	DCO	LOOP CONTROL	MODE
0	0	0	0	ON	ON	Active
0	0	0	1	ON	ON	LPM0
0	1	0	1	ON	OFF	LPM1
1	0	0	1	OFF	OFF	LPM2
1	1	0	1	OFF	OFF	LPM3
x	X	1	1	OFF	OFF	LPM4

2.9.2 System-Clock Frequency Control, SCFQCTL, address 052h

Bit 7: M. Modulation enable: enabled when M=0, disabled when M=1
Bit 6: N. Frequency multiplication factor: MCLK = N x ACLK

Where N = SCFQCTL(6..0) + 1.

2.9.3 System Clock Frequency Integrator 0, SCFI0, address 050h

2.9.4 System Clock Frequency Integrator 1, SCFI1, address 051h

These two registers contain the FN_4, FN_3, FN_2, $N_{DCO,}$ and N_{DCOmod} bits described in the *Modulator and DCO Range Control* sections.

2.9.5 Interrupt Enable Register 1, IE1, address 000h

Bit 1. OFIE: oscillator fault interrupt enable

2.9.6 Interrupt Flag Register 1, IFG1, address 002h

Bit 1. OFIFG: oscillator fault interrupt flag

2.10 References to Architecture Guide and Application Report

The references listed here can be found in the following documents:

AGML—*Architecture Guide and Module Library*, literature number SLAUE10B, 1996.

AR—Application report, literature number SLAAE10C, 1998

Interrupt processing	AGML	3–8
Interrupt related special function registers IEx, IFGx	AGML	3–10
Interrupt vector addresses	AGML	3–13
Operating modes	AGML	3–16
Low power modes	AGML	3–19
Oscillator and system clock generator	AGML	7–1
System clock control registers	AGML	7–9
DCO characteristics	AGML	7–12
Crystal oscillator module	AGML	16–3
Operating modes used on MSP430 Applications	AR	1–11
The system clock generator	AR	6–192
Accuracy of the system clock generator	AR	6–196
Design checklist	AR	7–10
Common software errors	AR	7–11
Common hardware errors	AR	7–14
Problems checklist	AR	7–15

IMPORTANT NOTICE

Texas Instruments and its subsidiaries (TI) reserve the right to make changes to their products or to discontinue any product or service without notice, and advise customers to obtain the latest version of relevant information to verify, before placing orders, that information being relied on is current and complete. All products are sold subject to the terms and conditions of sale supplied at the time of order acknowledgement, including those pertaining to warranty, patent infringement, and limitation of liability.

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

CERTAIN APPLICATIONS USING SEMICONDUCTOR PRODUCTS MAY INVOLVE POTENTIAL RISKS OF DEATH, PERSONAL INJURY, OR SEVERE PROPERTY OR ENVIRONMENTAL DAMAGE ("CRITICAL APPLICATIONS"). TI SEMICONDUCTOR PRODUCTS ARE NOT DESIGNED, AUTHORIZED, OR WARRANTED TO BE SUITABLE FOR USE IN LIFE-SUPPORT DEVICES OR SYSTEMS OR OTHER CRITICAL APPLICATIONS. INCLUSION OF TI PRODUCTS IN SUCH APPLICATIONS IS UNDERSTOOD TO BE FULLY AT THE CUSTOMER'S RISK.

In order to minimize risks associated with the customer's applications, adequate design and operating safeguards must be provided by the customer to minimize inherent or procedural hazards.

TI assumes no liability for applications assistance or customer product design. TI does not warrant or represent that any license, either express or implied, is granted under any patent right, copyright, mask work right, or other intellectual property right of TI covering or relating to any combination, machine, or process in which such semiconductor products or services might be or are used. TI's publication of information regarding any third party's products or services does not constitute TI's approval, warranty or endorsement thereof.