

How to Set Up and Use the **ADuCM4050**

FEATURES

Up to 52 MHz ARM Cortex-M4F processor
512 kB of embedded flash memory with ECC
Optional 4 kB cache for lower active power
128 kB system SRAM with parity
Power management unit
Multilayer advanced microcontroller bus architecture bus matrix
Central DMA controller
Beeper interface
SPORT, SPI, I²C, and UART peripheral interfaces
Cryptographic hardware support with AES-128, AES-256 along ECB mode, CBC mode, CTR mode, CCM mode, modified CCM mode, and SHA-256
Protected key storage with key wrap and unwrap
Keyed HMAC with key unwrap
2 RTCs
3 general-purpose timers
1 watchdog timer
RGB timer for driving RGB LED
Programmable GPIO pins
Hardware CRC calculator with programmable generator polynomial
Power-on reset and power supply monitor
12-bit successive approximation register ADC
True random number generator

GENERAL DESCRIPTION

This user guide provides detailed information on the [ADuCM4050](#) microcontroller functionality and features. Each section describes a different feature.

The [ADuCM4050](#) processor is an ultra low power, integrated, mixed-signal, microcontroller system used for processing, control, and connectivity. The microcontroller unit (MCU) subsystem is based on the ARM® Cortex-M4F processor, a collection of digital peripherals, cache embedded static random access memory (SRAM) and flash memory, and an analog subsystem that provides clocking, reset, and power management capabilities along with the analog-to-digital converter (ADC).

The [ADuCM4050](#) processor provides a collection of power modes and features, such as dynamic and software controlled clock gating and power gating, to support extremely low dynamic and hibernate power management.

Full specifications on the [ADuCM4050](#) are available in the product data sheet and the [ADuCM4050 Ultra Low Power ARM Cortex-M4F MCU with Integrated Power Management Hardware Reference Manual](#).

TABLE OF CONTENTS

Features	1	Effects of Cache	28
General Description	1	Current Consumption Comparison	30
Revision History	2	Dual RTC Feature in the ADuCM4050.....	31
Getting Started	3	Comparison of the RTC Features.....	31
Software Installation	4	Power Considerations	31
IAR Embedded Workbench Configuration	5	Conclusion	31
Power Optimization for the ADuCM4050 Processor	7	Benefits of ADuCM4050 the DC-to-DC Converter	32
ADuCM4050 Processor Power Management.....	7	DC-to-DC Basics.....	32
ADuCM4050 Processor Power Modes.....	7	Capacitors vs. Inductor Converters	34
Fast Wake Up from Hibernate Mode	13	Conclusions.....	35
Flash Memory and Instruction SRAM.....	13	UART Software Flow Control.....	36
Normal Wake Up.....	13	UART Flow Control	36
Fast Wake Up	13	System Description	37
Using the ADuCM4050 Processor Boot Kernel.....	17	Data Capture.....	40
Device Information Space Overview.....	17	SPI Flow Control Methods.....	41
Boot Kernel Overview	17	SPI Read Command Mode	41
UART Downloader	20	Flow Control Modes	43
Read Protection Key and Hashing	24	Conclusions	44
Memory Configuration	25	Sleep on Exit.....	46
Handling CRC in the IAR Workbench.....	26	Benefits	46
CrossCore Serial Flash Programmer	26	Enabling the Sleep on Exit Feature	46
Cache Memory in the ADuCM4050	28	System Control Register in the ADuCM4050.....	47
Storage and Access Ecosystem Block Diagram	28		
Flash Controller	28		

REVISION HISTORY

6/2018—Revision A: Initial Version

GETTING STARTED

This section introduces the tools and support packages required to develop an application for the ADuCM4050 microcontroller. This section describes how to download, install, and configure the files that program the ADuCM4050.

This section describes different steps in developing an application by using the IAR Embedded Workbench® as an integrated development environment (IDE). This section also describes how to download and run sample codes provided with the device family package (DFP) drivers.




Figure 1. IAR Embedded Workbench

SOFTWARE INSTALLATION

The software tools required to develop applications with the ADuCM4050 are available for download on the [EV-COG-AD4050](#) page.

Table 1. Required Software Tools

Tool	Functions
IAR Embedded Workbench	Used for compiling, debugging, and code development
Cross Core Embedded Studio	Used for compiling, debugging, and code development
Segger J-Link Software	J-Link software and documentation pack that includes USB drivers for the emulator, J-Link commander, and so on
ADuCM4050 Device Family Pack Drivers	Includes ADuCM4050 peripheral drivers and libraries, IAR configuration files, and the source and header files

Installing the Segger J-Link Driver

The Segger J-Link USB driver must be installed before using a serial wire interface, such as the interface of the IAR Embedded Workbench, to download and debug code.

Use the following procedure to install the J-Link USB driver:

1. Download the latest Segger J-Link software and documentation pack from the Segger website.
2. Run the executable software installer within the download directory.
3. Follow the on-screen instructions to complete the installation. Ensure the **Install USB Driver for J-Link** option is checked, as shown in Figure 2.




Figure 2. Segger J-Link Driver Installation Options

4. Plug in the J-Link emulator board, and open the **Device Manager**.
5. Check that the emulator board appears in the Windows® **Device Manager** in the USB controllers lists (see Figure 3).




Figure 3. Device Manager

After following the software installation procedures, the USB driver for the J-Link is installed and verified.

IAR Tools Installation

The IAR Embedded Workbench and the included IAR C/C++ Compiler generates the fastest performing, most compact code in the industry for ARM-based applications. Therefore, Analog Devices, Inc., created the DFP drivers for the ADuCM4050 for the IAR Embedded Workbench.

The KickStart edition is a free starter kit and evaluation version of IAR. This edition has limitations, both in code size (32 kB) and in the service and support provided.

The IAR KickStart software is available as a free trial to download on the IAR website.

For a detailed procedure on installing the IAR Embedded Workbench and adding license details, if required, refer to the *Installation and Licensing Guide for IAR Embedded Workbench*, which is available on the IAR website.

The DFP examples are for the IAR Embedded Workbench Version 8.20.1 or later. Project compatibility issues may occur when using different versions.

ADuCM4050 Device Family Pack (DFP)

The ADuCM4050 DFP provides the configuration, support files, and components required to ease the development of the ADuCM4050.

The contents of the DFP are as follows:

- Source files for the device drivers and services for use on the ADuCM4050 processor
- Tool chain support. These components are installed in the IAR Embedded Workbench workspace to configure the tool chain to recognize the ADuCM4050
- Documentation containing details about the present version of DFP

The IAR Embedded Workbench must be installed before installing the DFP.

Use the following procedure to install the DFP:

1. Download the ADuCM4050 DFP on the [EV-COG-AD4050](#) page.
2. Select the **ADuCM4x50 Device Family Pack**, which leads you to the page to download the latest version of the DFP. This is a CMSIS pack file that must be installed in IAR Embedded Workbench Version 8.11.3 and above.
3. Open the **IAR Embedded Workbench IDE** and select the **CMSIS Pack Installer**, as shown in Figure 4.



Figure 4. CMSIS Pack Installer

4. Click **Install local pack file** in the window that is shown in Figure 5, and browse to the DFP CMSIS pack file downloaded to install it.




Figure 5. Install Local Pack File Button

5. After installation completes, the DFP displays installed (see Figure 6).



Figure 6. Completed Installation of the DFP

IAR EMBEDDED WORKBENCH CONFIGURATION

This section describes the IAR Embedded Workbench configuration procedure for proper operation of the ADuCM4050. Only the sections that must be modified from the default values are described.

Take the following steps to configure the IAR:

1. In **Options**, ensure the **Analog Devices ADuCM4050** device option is selected as the target.

2. Go to **C/C++ Compiler > Optimizations**, select from the different optimization options for speed, code size, balance, and so on, depending on your application needs. Sometimes, the compiler identifies the writes of a register as eligible to be optimized, which may cause unexpected behavior. In such situations, it is recommended to protect configuration functions from being optimized by using the following code:

```
#pragma optimize=none
```
3. Go to **C/C++ Compiler > Preprocessor** and include the path of the included directories, depending on the code to be run.
4. Go to **Debugger** and select **CMSIS DAP** as the debugger within the **Driver** dropdown menu (see Figure 7). Verify that both the **Verify download** and **Use flash loader(s)** boxes are checked in the **Debugger > Download** menu, as shown in Figure 8.




Figure 7. Debugger > Setup Tab




Figure 8. Debugger > Download Tab

5. Figure 9 shows the **CMSIS DAP > Setup** configuration. Be sure to use the **Hardware** target reset strategy under the **Reset** dropdown menu. Figure 10 shows the **CMSIS DAP > Interface** configuration




Figure 9. CMSIS DAP > Setup Configuration Tab




Figure 10. CMSIS DAP > Interface Configuration Tab

POWER OPTIMIZATION FOR THE ADuCM4050 PROCESSOR

Choosing a low power MCU is a difficult task because it involves reviewing data sheets to analyze electrical specifications. It is often difficult to relate these specifications to applicable, system level use cases.

Evaluating various power modes while considering peripheral operations emulating real use case scenarios is an essential step in choosing the right MCU for a power sensitive application. Key aspects to evaluate when choosing an MCU for low power applications include the following:

- Availability of low power modes and the impact of these modes on the ability to retain the contents of the SRAM.
- Power consumption with the real-time clock (RTC) running while the rest of the system is in low power mode.
- Wake-up times from low power modes.
- Supply voltage range from an application standpoint. The designer can adjust and select the system supply voltage, depending on the component requirements.
- Power consumption in active mode.
- Core activity—example algorithm processing.
- Peripheral activity—direct memory access (DMA) operations.
- Simultaneous core and peripheral activity.
- Flexibility in choosing core and peripheral clock frequencies that meet system requirements while keeping the power consumption low.
- Hardware DMA blocks that enable the CPU to be in low power mode during peripheral activity.

The ADuCM4050 processor is an ultra low power, integrated, mixed-signal, MCU system for processing, control, and connectivity. The MCU system is based on an ARM Cortex-M4F processor, offering 1.25 DMIPS per MHz of performance running up to 52 MHz, combined with a collection of digital peripherals, embedded SRAM and flash memories, and an analog subsystem that provides clocking, reset, and power management capabilities in addition to an ADC subsystem.

The ADuCM4050 processor is one of the few ultra low power MCUs on the market that offers a cache controller. Programs that repeatedly access the same data or instructions make

effective use of cache memory, thereby reducing overall power consumption.

The power consumption of an MCU largely depends on two factors: the operating voltage and the frequency at which the system operates. The ADuCM4050 processor incorporates several power modes that are useful in building battery-powered or self-powered (energy harvesting) applications.

This section discusses the power modes of the ADuCM4050 processors in detail and provides example power measurements for several scenarios, to help developers choose power modes that best fit low power application requirements.

ADUCM4050 PROCESSOR POWER MANAGEMENT

The ADuCM4050 processor incorporates a highly customizable power management and clocking system that offers application developers the flexibility to balance power and performance. The power management blocks consist of integrated regulators, a clock gating scheme, and switches applicable to numerous application scenarios.

The power management system features are as follows:

- An integrated 1.2 V low dropout (LDO) regulator and an optional capacitive buck regulator.
- Integrated power switches for low standby current in hibernate mode.
- Power gating to reduce leakage in sleep modes.
- A power supply monitor with a selectable voltage range.

ADUCM4050 PROCESSOR POWER MODES

The power management system provides the following low power modes:

- Active mode with customized clock gating features.
- Flexi mode with smart peripherals.
- Hibernate mode with optional SRAM retention capability.
- Shutdown mode without SRAM retention.
- Shutdown fast mode without SRAM retention.

Each mode provides a low power benefit with potential functionality trade-offs.

Table 2 summarizes the status of the system blocks in each low power mode.

Table 2. Power Mode System Block Status

Functional Block	ARM Cortex-M4F Core	Buck	Peripheral/DMA	HFXTAL	HFOSC	LFXTAL	PLL	LFOSC	RTC0	RTC1	ADC	SRAM	Flash
Active Mode	On	User ¹	User ¹	User ¹	User ¹	User ¹	User ¹	On	User ¹	User ¹	User ¹	On	On
Flexi Mode	Off	User ¹	User ¹	User ¹	User ¹	User ¹	User ¹	On	User ¹	User ¹	User ¹	On	On
Hibernate Mode	Off	Off	Off	Off	Off	User	Off	On	User ¹	User ¹	Off	On ²	Off
Shutdown Mode	Off	Off	Off	Off	Off	User	Off	Off	User ¹	Off	Off	Off	Off
Shutdown Fast Mode	Off	Off	Off	Off	Off	User	Off	Off	User ¹	Off	Off	Off	Off

¹ In the user application code, this functional block can be configured on or off.

² The retainable SRAM size is configurable.

Active Mode

In active mode (also called full on mode), the ARM Cortex-M4F is active and executes instructions from flash memory and/or SRAM. All peripherals can be enabled or disabled at the discretion of the user, and active mode power can be enhanced by optimized clock management.

Several power saving options are available in active mode:

- Using the buck converter.
- Enabling the cache.
- Using dynamic clock scaling.
- Using clock gating.

Buck Converter

The optional integrated buck converter feature saves power in active mode. The buck converter powers the linear regulator, which powers the digital core domain. The buck converter enters bypass mode after the battery voltage (V_{BAT}) falls below ~2.3 V. After entering bypass mode, the buck converter output follows the input.

For designs in which the optional buck converter is not used, the VDCDC_CAP1P, VDCDC_CAP1N, VDCDC_OUT, VDCDC_CAP2P, and VDCDC_CAP2N pins must be left unconnected.

The buck converter is solely for processor usage. An external load cannot be connected to the buck converter output.

Enable the buck converter by setting the HPBUCKEN bit in the CTL1 register per the following code:

```
*pREG_PMG0_CTL1 |= (1 << BITP_PMG_CTL1_HPBUCKEN);
```

Figure 11 compares the power consumption of the ADuCM4050 processor when computing prime numbers with the following conditions:

- $V_{BAT} = 3.0$ V
- $HCLK = PCLK = 26$ MHz
- Cache memory disabled

The buck converter impacts current consumption positively at higher V_{BAT} values. Specifically, there is roughly a 50% decrease in the active current when $V_{BAT} \geq 3$ V.




Figure 11. Impact of the Buck Converter on Active Mode Power Consumption

Enabling Cache Memory

Cache memory reduces the average time to access data from flash memory. For scenarios in which the CPU is required to run an algorithm or the same data must be accessed repeatedly, cacheable memory can reduce the power consumption because execution is from the internal instruction SRAM. When the cache controller is enabled, 4 kB of instruction SRAM is reserved as cache memory.

Cache memory is disabled at startup by default. Use the following procedure to enable the cache memory:

1. Read the cache enable status bit (Bit 0 in the FLCC_CACHE_STAT register) to ensure cache memory is disabled. Poll this bit until it clears.
2. Write the user key to the FLCC0_CACHE_KEY register.

```
*pREG_FLCC0_CACHE_KEY = 0xF123F456;
```
3. Set the instruction cache enable bit (ICEN in the FLCC0_CACHE_SETUP register) as follows:

```
*pREG_FLCC0_CACHE_SETUP |= (1 << BITP_FLCC_CACHE_SETUP_ICEN);
```

Figure 12 compares the power consumption of the ADuCM4050 processors when computing prime numbers with the following conditions:

- $V_{BAT} = 3.0\text{ V}$
- $HCLK = PCLK = 26\text{ MHz}$
- Buck converter disabled




Figure 12. Impact of Cache Memory on Active Mode Power Consumption

Enabling the cache memory reduces the average active current consumption by ~15%.

Dynamic Clock Scaling

Dynamic clock and/or frequency scaling is a proven method to reduce power consumption. The ADuCM4050 processors have a flexible clock architecture that allows dynamic modification of the CPU and peripheral clock frequencies. A combination of clock dividers and a phase-locked loop (PLL) provides flexibility in deriving an optimum system clock frequency that guarantees system performance while keeping the power consumption low, as compared to a fixed clock scheme. Programmable clock dividers are available to generate the clocks in the system, and the divisors can be configured while the code is running.

Figure 13 plots the power consumption of the ADuCM4050 processors when computing prime numbers with the following conditions:

- $V_{BAT} = 3.0\text{ V}$
- $HCLK = PCLK$ (the source of the root clock is HFOSC)
- Buck converter disabled
- Cache disabled




Figure 13. Impact of Core Clock Frequency on Active Mode Power Consumption

Figure 14 plots the power consumption of the ADuCM4050 processors when computing prime numbers with the following conditions:

- $V_{BAT} = 3.0\text{ V}$
- $HCLK = PCLK$ (the source of the root clock is XTAL)
- Two wait state on the flash
- Buck converter disabled
- Cache disabled

Power dissipation decreases as core clock frequency decreases.




Figure 14. Impact of Core Clock Frequency on Active Mode Power Consumption with the PLL

Clock Gating

The system is heavily clock gated and uses automatic clock gating techniques. Most peripherals are automatically clock gated when the peripheral is disabled, such that the clock runs only when the peripheral is enabled. The exceptions are I²C, general-purpose input/output (GPIO), and the general-purpose timer (GPTMR). These blocks must be manually clock gated using the CLKCON5 register. Gate the peripheral clock completely by setting the CLKG_CLK_CTL5 bit. Any access to the clock gated peripherals overrides the clock gate settings in the CLKCON5 register.

For application scenarios in which the core is processing data and no peripheral activity is desired, the peripheral clock (PCLK) can be turned off to save power. Figure 15 shows the power consumption of the ADuCM4050 processors when computing prime numbers with the following conditions:

- $V_{BAT} = 3.0\text{ V}$
- Buck converter disabled
- Cache memory disabled
- HCLK = PCLK = 26 MHz




Figure 15. Impact of Peripheral Clock Gating on Active Mode Power Consumption

As shown in Figure 15, a ~0.2 mA reduction in the active current is observed when the peripheral clock is gated.

In active mode, the four techniques described in this section can be combined to achieve maximum power savings.

Flexi Mode

Flexi mode is a flexible sleep mode useful in scenarios in which the core must wait for a peripheral data transfer to complete before it can start processing. In flexi mode, the core is clock gated, and the remainder of the system is active. Flexi mode substantially reduces active power when low speed activity is expected to complete (for example, reading a certain number of bytes from a sensor) before the processor must be woken up to process the data.

Consider a scenario in which the CPU configures a serial peripheral interface (SPI) DMA and must wait for the DMA to complete. Figure 16 shows the power consumption of the ADuCM4050 processors transferring data over the SPI using DMA accesses with the following conditions:

- $V_{BAT} = 3.0\text{ V}$
- Buck converter enabled
- Cache disabled
- PCLK = 26 MHz
- SPI_DIV = 49




Figure 16. Impact of Flexi Mode on Power Consumption

As shown in Figure 16, there is nearly a 71% savings in power when flexi mode is used while the DMA is ongoing, rather than keeping the core in active mode.

A number of wake-up sources that exit flexi mode (for example, DMA interrupts, external interrupts, timer interrupts, and so on), and it typically takes only one CPU clock cycle to exit.

The buck converter can also be enabled in flexi mode to save additional power. Figure 17 shows the power consumption of the ADuCM4050 processor across V_{BAT} in flexi mode with the buck converter on while transferring data over the SPI using DMA accesses with the following conditions:

- $V_{BAT} = 3.0\text{ V}$
- Cache memory disabled
- PCLK = 26 MHz
- SPI_DIV = 49




Figure 17. Impact of the Buck Converter on Flexi Mode Power Consumption

As seen in Figure 17, a similar power improvement pattern to the impact of the buck converter in active mode is shown in Figure 11. Specifically, when $V_{BAT} \geq 3\text{ V}$, a 50% improvement in power is observed.

Hibernate Mode

In hibernate mode, the ARM Cortex-M4F core and all digital peripherals are off with configurable SRAM retention, port pin retention, a limited number of wake-up interrupts, and, optionally, an active RTC. All GPIO pin states are retained in hibernate mode. The ADuCM4050 processor also incorporates the SensorStrobe™ mechanism in the RTC1 block, which enables ultra low power sensor data measurement.

Before entering hibernate mode, most of the enabled peripherals must be programmed to undergo a specific sequence to properly enter or exit hibernate mode, and several system memory map registers (MMRs) and peripheral registers are retained while in hibernate mode. For more details, refer to the relevant peripheral information in the [ADuCM4050 Ultra Low Power ARM Cortex-M4F MCU with Integrated Power Management Hardware Reference Manual](#).

Configurable Retainable SRAM

The ADuCM4050 processor supports SRAM partitioned in 32 kB blocks. Several memory sizes can be retained from 16 kB (default) up to 124 kB while in hibernate mode. The more SRAM that must be retained, the higher the power consumption in hibernate mode, as shown in Figure 18.




Figure 18. Current on the V_{BAT} Supply Pin (I_{BAT}) for Various Retained SRAM Values

The SRAM retention size can be configured by setting the appropriate bits in the PMG0_SRAMRET register. For example, to enable 124 kB of SRAM to be retained while in hibernate mode, use the following code in the PMG0_SRAMRET register:

```
*pREG_PMG_PWRKEY = 0x4859;
*pREG_PMG_SRAMRET |=
((1 << BITP_PMG_SRAMRET_SRAM_RET1_EN) |
(1 << BITP_PMG_SRAMRET_SRAM_RET2_EN) |
(1 << BITP_PMG_SRAMRET_SRAM_RET3_EN) |
(1 << BITP_PMG_SRAMRET_SRAM_RET4_EN));
```

If parity is enabled, initialization of nonretained SRAM regions may be required upon waking from hibernate mode.

Wake-Up Sources

The following events are capable of waking the device up from hibernate mode:

- External Interrupt 0 to External Interrupt 3
- RTC0 and RTC1 interrupt
- Battery voltage range interrupt
- Universal asynchronous receiver/transmitter (UART) receiver (Rx) pin activity

Of the two real-time clocks, RTC1 is the recommended wake-up source from hibernate mode because it consumes less power than RTC0, and it provides more functionality. In an application where both hibernate mode and shutdown mode must be used, only RTC0 can be used. RTC0 is the only real-time clock that can be used for exiting shutdown mode. See the Dual RTC Feature in the ADuCM4050 section for more information.

The wake-up time from hibernate mode from any of these events is ~10 μs when executing from flash and ~5 μs when executing from SRAM.

RTC Clock Sources

The ADuCM4050 processor offers two clock choices for the RTC1 block:

- A low power internal RC oscillator (LFOSC).
- An external crystal oscillator (LFXTAL).

Consider that the input clock for RTC1 is LFMUX. Therefore, the decision of choosing between LFOSC and LFXTAL must be made while taking into account the requirements for the other peripherals clocked by LFMUX.

Choosing to implement either LFOSC or LFXTAL is a trade-off between accuracy and power consumption. LFXTAL is more accurate (depending on the crystal manufacturer) compared to the LFOSC, but LFOSC dissipates less power, as shown in Figure 19.




Figure 19. Hibernate Current with RTC1 as the Wake-Up Source (LFOSC vs. LFXTAL)

SensorStrobe

The SensorStrobe mechanism allows the ADuCM4050 processor to be used as a programmable clock generator in all power modes, except in shutdown and shutdown fast modes. In this way, the external sensors have their timing domains mastered by the ADuCM4050 processors because the SensorStrobe output is a programmable divider from RTC1, which can operate up to a resolution of 30.7 μ s. The sensors and microcontroller are synchronous, which removes the need for additional resampling of data to time align the microcontrollers and the sensors.

Shutdown Mode

Shutdown mode is the deepest sleep mode in which digital and analog circuits are powered down. The state of the digital core and the SRAM memory content is not retained. However, the state of the pads is preserved, as is the wake-up interrupt configuration.

The configuration of the pads is preserved and locked after waking up from shutdown mode. The user must unlock the state of the pads by writing 0x58FA to the PGM_TST_CLR_LATCH_GPIO register, preferably inside the interrupt service routine (ISR) routine:

```
*pREG_PMG0_TST_CLR_LATCH_GPIO = 0x58FA;
```

Additionally, the user must configure the appropriate wake-up source, choosing from the following options:

- External Interrupt 0 to External Interrupt 2
- External reset
- Battery falling below 1.6 V
- RTC0 timer

The RTC0 block can optionally be enabled in this mode, which allows the processor to be periodically woken up by the RTC0 interrupt.

The clock source for RTC0 must be LFX TAL because LFOSC is disabled in shutdown mode.

Because the RTC0 block must be powered to serve as a wake-up source, it adds to the power dissipation while in shutdown mode, as shown in Figure 20.

When the device wakes up from shutdown mode, the power-on reset (POR) sequence is followed, and code execution starts from the beginning.




Figure 20. Shutdown Mode Current (External Sources vs. RTC0)

Shutdown Mode—Fast Wake Up

Fast wake up mode has exactly the same behavior as the mode described in the Shutdown Mode section, except for the faster wake-up time at the expense of higher power consumption.

Figure 21 shows the power consumption in this power mode. As shown, enabling the RTC0 adds extra power dissipation.




Figure 21. Shutdown Mode Current (External Sources vs. RTC0)

FAST WAKE UP FROM HIBERNATE MODE

FLASH MEMORY AND INSTRUCTION SRAM

Flash memory is the long-term storage medium for any microcontroller. Because the flash memory is nonvolatile, it is used for the storage of constant data and program code of the microcontroller. Due to the nature of flash memories, memory access is slow compared to SRAM, cache, and other registers. The latency takes effect during the execution of a looping code that performs real-time calculations.

The ADuCM4050 microcontroller has the capacity to execute instructions in SRAM. Instruction SRAM (iSRAM) is a portion of the ADuCM4050 microcontroller dedicated to be a temporary program code and instruction storage. It is used when the program must execute a looping code faster and does not want the flash latency to affect the execution. The iSRAM is small (up to 32 kB) compared to the 512 kB flash memory. Therefore, only important code and instructions must use the iSRAM.

NORMAL WAKE UP

Cortex-M4F core and all the digital peripherals (except some user-selectable SRAM blocks) are turned off during hibernate mode. The RTC can be configured by the user program to turn off during hibernate. (RTC0 and/or RTC1 might be on depending on the user configuration). Registers of the digital peripherals are also turned off. This is to ensure the low current consumption of the microcontroller during hibernate mode, although some registers are retained to allow the device to wake up in the same status it was in when it went to sleep.

An interrupt from one of the allowed wake-up sources boots the microcontroller from hibernate to active mode to service the interrupt. Refer to the [ADuCM4050 Ultra Low Power ARM Cortex-M4F MCU with Integrated Power Management Hardware Reference Manual](#) document for further information about the possible interrupt wake-up sources. During the transition from hibernate to active mode, the microcontroller reinitializes the digital peripherals that are off during hibernate before it executes the first instruction on the ISR.

The flash memory, where the program code is located, is also off during hibernate mode. Upon microcontroller wake up, the core turns on the flash memory. The flash memory initialization is slow, and it takes about 5.7 μ s to complete before it can take commands from the controller.

Figure 22 shows the delay time from triggering an external wake-up signal to microcontroller response by lighting up a light emitting diode. The red trace is the signal from the push button that triggers an external interrupt signal to wake up the microcontroller. The blue trace is the GPIO toggling indicating that the microcontroller is awake, that is, the first instruction executed is a GPIO toggle (GPIO_TGL). The entire wake-up process takes around 10 μ s to complete.




Figure 22. Normal Wake Up from Hibernate Mode Response

FAST WAKE UP

One way to wake up from hibernate faster is to move the required functions and instructions to iSRAM. The microcontroller code starts earlier because the code must not wait for the completion of the flash memory initialization.

Figure 23 shows the delay time from triggering an external wake-up signal to microcontroller response with the fast wake-up procedure. The red trace is the signal from the push button, and the blue trace is the GPIO toggled as first instruction after waking up.

The wake-up time is reduced to around 4.5 μ s. This time assumes a 50% improvement with respect to normal wake up.




Figure 23. Fast Wake Up from Hibernate Mode Response

Table 3. Wake-Up Times with Different Peripherals

ADuCM4050 Peripheral Activity	Normal Wake-Up Time (μs)	Fast Wake-Up Time (μs)
Pin Toggle (Lighting an LED)	9.7	4.4
I ² C Clock (First Edge, 400 kHz)	12.3	7.0
SPI Clock (First Edge, 400 kHz)	12.5	7.2

Table 3 shows a comparison of the wake-up times to different peripherals. The wake-up time is measured from the falling edge of the external wake-up signal to the first rising edge of the serial clocks (as for I²C and SPI) or the rising edge of the pin connected to the light emitting diode.

To perform fast wake up from hibernate mode, use the following procedure:

1. Initialize the iSRAM.
2. Modify the linker script to add sections for the functions and to remap the addresses.
3. Relocate the interrupt vector table (VTOR) from flash to SRAM.
4. Place the required functions and interrupt handlers in iSRAM.

Initialization of the iSRAM

To use the iSRAM, take the following steps:

1. Enable the iSRAM bank by asserting the INSTREN bit field in the PMG_TST_SRAM_CTL register.
2. Retain the half of the iSRAM by asserting the BNK2EN bit in the PMG_TST_SRAM_CTL register.

Modification of the Linker Script

Modify the linker script to help the linker place the code in the correct place in the memory map. Remap the addresses, and add the required sections for the location of the program code and interrupt handlers to iSRAM.

If using the IAR Embedded Workbench, apply the following changes to the linker script:

1. Remap the SRAM addresses. Search for the following line in the linker script:

```
define symbol USER_SRAM_MODE = 0;
```

Change the value of the USER_SRAM_MODE to 0 or 1.

2. Add sections for iSRAM. Search the linker script for the following lines:

```
// iSRAM section for placing code in SRAM
place in iRAM {section ISRAM_REGION };
initialize by copy {section ISRAM_REGION
};
```

Modify the lines to include the .textw section. IAR linker places the instructions for SRAM in the .textw section.

```
// ISRAM section for placing code in SRAM
place in iRAM {section ISRAM_REGION,
section .textw};
```

```
initialize by copy {section ISRAM_REGION,
section .textw};
```

Interrupt Vector Table Relocation

The interrupt vector table (IVT) lists the different interrupt sources for the ADuCM4050 microcontroller. The following events are capable of waking the MCU up from hibernate mode:

- External Interrupt 0 to External Interrupt 3
- RTC0 and RTC1 interrupt
- Battery voltage range interrupt
- UART receiver (Rx) pin activity

For fast wake up, move the IVT from flash to SRAM. In this way, the microcontroller does not need to wait for the flash initialization to check the location of the interrupt handler of the wake-up source.

Copy the IVT to SRAM and update the SCB register to the VTOR of the address of the interrupt vector in SRAM.

The file system ADuCM4050.c must be modified as follows:

1. Add multiple definitions before the SystemInit function

```
#ifndef RELOCATE_IVT
#include <assert.h>
#ifdef __GNUC__
#define ATTRIBUTE_INTERRUPT
__attribute__((interrupt__))
#define KEEP_VAR(var) var
__attribute__((used))
#define WEAK_PROTO(proto)
__attribute__((weak)) proto
#define WEAK_FUNC(func)
__attribute__((weak)) func
#define VECTOR_SECTION ".isr_vector"
#define SECTION_PLACE(def, sectionname)
__attribute__((section(sectionname))) def
#define RESET_EXCPT_HNDLR ResetISR
#define COMPILER_NAME "GNUCC"
#endif // __GNUC__

#ifdef USER_SPECIFIED_RTOS
#include <user_rtos_support.h>
#endif
```

```
#ifndef __ARMCC_VERSION
#define ATTRIBUTE_INTERRUPT
```

```

#define KEEP_VAR(var) var
__attribute__((used))
#define WEAK_PROTO(proto) proto
__attribute__((weak))
#define WEAK_FUNC(func) func
#define VECTOR_SECTION "RESET"
#define SECTION_PLACE(def,sectionname)
__attribute__((section(sectionname))) def
#define RESET_EXCPT_HNDLR
__main
#define COMPILER_NAME
"ARMCC"
#endif // __ARMCC_VERSION

#ifdef __ICCARM__

/*
 * Pm154 (rule 19.10): in the definition of a
 * function-like macro, each instance of a
 * parameter shall be enclosed in parentheses.
 * The parameters in the following macros
 * cannot be enclosed in parentheses.
 */
#pragma diag_suppress=Pm154

#define ATTRIBUTE_INTERRUPT
#define KEEP_VAR(var) __root var
#define WEAK_PROTO(proto) __weak proto
#define WEAK_FUNC(func) __weak func
#define VECTOR_SECTION ".intvec"
#define SECTION_PLACE(def,sectionname) def @
sectionname
#define RESET_EXCPT_HNDLR
__iar_program_start
#define COMPILER_NAME
"ICCARM"

#endif /* __ICCARM__ */

extern uint32_t __Vectors_Size;
#define RELOCATION_ADDRESS (0x20000000)
#define RELOCATION_ALIGNMENT (0x200)

#define LENGTHOF_IVT (88u)

#if defined ( __ICCARM__ )
    #pragma
    data_alignment=RELOCATION_ALIGNMENT /* IAR
    */
#elif defined (__CC_ARM)

```

```

    __align(RELOCATION_ALIGNMENT)
    /* Keil */
#else
    #pragma message("WARNING: NO ALIGNMENT
    DEFINED FOR IVT RELOCATION")
#endif
/* reserve no-init aligned IVT space at top
of RAM */
SECTION_PLACE(KEEP_VAR(__no_init void *
__relocated_vector_table[LENGTHOF_IVT]),
RELOCATION_ADDRESS);
#endif

```

2. Replace this line

```
SCB->VTOR = (uint32_t) &__Vectors;
```

with these lines

```

#ifdef RELOCATE_IVT
    /* Copy the IVT (avoid use of memcpy
    here so it does not become locked into
    flash). */
    size_t i;
    //assert(
    sizeof(__relocated_vector_table) ==
    __Vectors_Size);
    for (i = 0u; i < LENGTHOF_IVT; i++) {
        __relocated_vector_table[i] =
        __Vectors[i];
    }
    SCB->VTOR = (uint32_t)
    &__relocated_vector_table;
#else
    /* Set the vector table address */
    SCB->VTOR = (uint32_t) &__Vectors;
#endif

```

3. Replace this line

```
extern uint32_t __Vectors;
```

with these lines

```

#ifdef RELOCATE_IVT
extern void* __Vectors[];
#else
extern uint32_t __Vectors;
#endif

```

Finally, it is necessary to declare the define RELOCATE_IVT by adding it to the **C/C++ Compiler/Preprocessor** tab (see Figure 24) to activate the relocation code built in the evaluation board support package.




Figure 24. IAR C/C++ Compiler Options—RELOCATE_IVT Directive in the Defined Symbols Textbox

Placement of Program Code to SRAM

The procedure of placing the program code to SRAM is dependent on the compiler used.

For most compilers, use the following line before the function definition to tell the location of the defined function:

`#pragma location="<linker_section>"`

For the IAR Embedded Workbench, use the `__ramfunc` keyword before the function (see Figure 25).

```
__ramfunc void Ext_Int0_Handler(void) {
    __disable_irq();

    // write to spi tx registers
    *((volatile uint32_t *) REG_SPI2_CTL) = spi_con_buff;
    *((volatile uint32_t *) REG_SPI2_TX) = 0x48;
    *((volatile uint32_t *) REG_SPI2_TX) = 0x88;

    *((volatile uint32_t *) REG_GPIO0_OUT) = (1 << 13);
    *((volatile uint32_t *) REG_XINT0_CLR) |= 1;

    SCB->SCR &= ~(1u << 1);
}
```

Figure 25. Sample for Appending `__ramfunc` Directive to a Function

The amount of iSRAM retained depends on the PMG0_SRAMRET register configuration. The RET1 bit field enables the retention of the 12 kB in the range of 0x1000_0000 to 0x1000_2FFF. The RET2 bit field enables the retention of the 16 kB in the range of 0x1000_3000 to 0x1000_6FFFF. Be sure that the location of the code is in the retained region.

USING THE ADuCM4050 PROCESSOR BOOT KERNEL

The ADuCM4050 processor features integrated flash memory that contains the user application code (user space) and a dedicated 4 kB bank of memory, with the information space arranged as shown in Figure 26. The total device storage is generally described as the size of the user space.




Figure 26. Flash Information Memory Space

The ADuCM4050 processor features 512 kB of user space.

As shown in Figure 26, the information space block is further broken down into the boot kernel, residing in the upper 4 kB of flash memory, and the device information space. The boot kernel is responsible for implementing a secure environment, where the user application code can optionally be read and/or write protected, and executes the application from flash memory on reset. The boot kernel also provides a mechanism to upgrade the firmware through a UART downloader.

This section describes the information space region of the on-chip flash memory, as well as both the boot process and how to use the UART downloader to perform field upgrades to the processor firmware.

DEVICE INFORMATION SPACE OVERVIEW

The device information space is reserved for use by Analog Devices and stores several trim and calibration values and other device specific metadata. The 256 bytes of device information space in Address 0x0008_0F00 to Address 0x0008_0FFF are protected and cannot be read by user code (attempted reads return bus error). The 32 bytes of device information space in Address 0x0008_0EE0 to Address 0x0008_0EFF can be read by user code and Table 4 summarizes the information stored therein. None of the device information space can be programmed or erased by the user.

Table 4. List of User Accessible Parameters in Device Information Space

Address Range	Size	Description
0x0008_0EF0 to 0x0008_0EFF	128 bits (16 bytes)	Unique ID
0x0008_0EE4 to 0x0008_0EEF	96 bits (12 bytes)	Manufacturer ID
0x0008_0EE0 to 0x0008_0EE3	32 bits (4 bytes)	Revision number of kernel

BOOT KERNEL OVERVIEW

The boot kernel switches to the user application after performing certain checks (including the CRC integrity of the user application), or the boot kernel enters UART downloader mode to upgrade the user application in flash memory, depending on the SYS_BMODE0 boot mode pin state at reset.

The boot kernel supports in field updates to the user application through the UART port. For security reasons, the boot kernel itself does not provide the flash programming feature. However, it allows the firmware update code, which has flash driver code for updating the user flash, to be downloaded to the device over the UART port. This code is referred to as a second stage loader (SSL) and is run from SRAM. The SSL must be authenticated before it can be provided run access. The security scheme implemented is discussed in the Read Protection Key Hash section and the Key Hash CRC section, describing the critical part of the kernel to provide the secure environment in which the user code can be read and/or write protected, allowing intellectual property security.

The serial download capability allows developers to reprogram the device while it is soldered directly onto the target system, avoiding the need for an external device programmer and removing the need to swap the device out of the system. The serial download feature also enables system upgrades to be performed in the field, provided the hardware infrastructure involving the SYS_BMODE0 pin and the UART port are implemented on the target board.

Configuring Security Options

The boot kernel provides the flexibility to configure the security options of the device by allowing the user to program certain keys and parameters in predefined locations in Page 0 of the user flash memory. The kernel provides the user code security and integrity, which depends on the number of user defined parameters in the first page of the user flash memory. Table 5 summarizes the list of keys and parameters, as well as their locations in the user flash memory.

Table 5. List of Keys and Parameters

Address Range	Size	Description
0x0000_0180 to 0x0000_018F	128 bits (16 bytes)	Read protection key hash
0x0000_0190 to 0x0000_0193	32 bits (4 bytes)	CRC of read protection key hash
0x0000_0194 to 0x0000_0197	32 bits (4 bytes)	Length of user boot loader or entire user code (used for CRC verification before boot)
0x0000_0198 to 0x0000_019B	32-bit word	In circuit write protection if set to no write (NOWR)
0x0000_019C to 0x0000_019F	32-bit word	CPU write protection of individual flash blocks

Read Protection Key Hash

Program the 128-bit read protection key hash at Address 0x00000180 in the first page of the user flash memory. The value of this hash depends on the kind of security desired in the system, because this security defines the read accessibility to the device. The key hash defines the state of the serial wire debugger (SWD), as well as the access permission of the SSL downloaded for upgrades via the UART.

The reset state of the flash memory of all logic high memory cells (along with a valid key hash CRC) indicates that the user does not desire read protection. In this case, the SWD interface is automatically enabled during booting.

Any nonreset value results in the SWD being locked. Therefore, there is no SWD access to the device.

The key hash is the 128-bit, truncated secure hash algorithm (SHA-256) of the user key (which is 128 bits in length), which can be sent along with the SSL during the UART download phase. If the user key is valid, and the hash of the received key matches the key hash stored, the SSL runs with all permissions. If the user key fails the key hash check, the SSL only has write permission to the user flash.

Key Hash CRC

The key hash has a 32-bit CRC checksum stored at Address 0x00000190. The key hash is valid only if the associated 4-byte checksum is valid. The key hash has a separate key hash to protect it against flash tempering attacks. The user must ensure a valid CRC for the key hash is stored along with the key hash itself.

In Circuit Write Protect Key

The 32-bit in circuit write protect key at Address 0x00000198 of the user flash memory prevents in circuit programming of the device. To disable in circuit reprogramming, program the hexadecimal value of the NoWr ASCII string (without the terminating null character) to this address. In this case, SWD access to the device is locked, and the only way to update the device code is via the UART downloader.

Use in circuit write protection along with read protection (providing both read and write protection). In circuit write protection alone does not have any significance.

Write Protection

Pages can be locked to prevent code from accidentally erasing and reprogramming critical flash memory blocks (such as the user code boot loader). There is a hardware register in the flash controller that disables the programming of pages grouped into blocks. This register is not automatically loaded via the hardware. Rather, this register is written to via the kernel. The kernel reads the write protection word from the user flash address, Address 0x0000019C, and writes it to the write protection register in the flash controller. The user can write the appropriate word to this location, depending on the pages intended to be protected against accidental writes. The pages are protected in groups of

four, with each bit in the 32-bit word corresponding to four continuous flash pages. Refer to the [ADuCM4050 Ultra Low Power ARM Cortex-M4F MCU with Integrated Power Management Hardware Reference Manual](#) for details.

User Code Length

There is a 32-bit value stored at a flash memory address, Address 0x00000194, that defines the CRC protected user code length. The value programmed in this field defines the page number of the user flash memory up to which the CRC protection is desired by the user. The value of N means that CRC protection is desired from Page 0 to Page N of the flash memory, protecting a total of N + 1 pages.

Valid values for the field in Address 0x00000194 are 0 to 255 for the 512 kB [ADuCM4050](#) processor. Any value outside this range is treated as invalid and results in a CRC check failure.

User Code CRC

The user code CRC is stored at the end of Page N. The 32-bit CRC (MSB first) with a polynomial of 0x4C11DB7 is expected by the kernel. If the page number is N, the CRC is expected to reside at flash memory address $(N \times 0x800) + 0x7FC$. For example, if N = 5, a total of six pages are CRC protected, and the CRC is stored at Address 0x2FFC. There is an option to disable the CRC check by programming 0xFFFFFFFF to the expected CRC location. After the kernel sees this value in the CRC location, it skips the CRC check.

Boot Code Flow

This section describes how the kernel operates, based on the user programmable parameters described in the Read Protection Key Hash section through the User Code CRC section. Figure 27 shows a flowchart of the boot code.

After reset, the boot kernel inspects all the parameters stored in Page 0 of the user flash memory. The user has not requested the read protection if the user read protection key hash is not programmed (meaning that it is set to all FFs) and the key hash CRC is valid. As such, the SWD is enabled. However, flash access may be protected, depending on the state of the user code CRC.

If the CRC is valid, access to user flash memory is unrestricted.

If the user disables the CRC by programming 0xFFFFFFFF to the CRC location, access to the user flash memory is also unrestricted.

If the CRC is invalid, the user flash memory is protected with no read or write access allowed. Only the flash mass erase command is allowed. In this case, user code execution is not allowed.

Take care to program a valid CRC (or 0xFFFFFFFF) in the defined CRC location. Otherwise, the user flash memory is read protected by the kernel. In this case, flash-based applications fail to load unless the user flash memory is mass erased.

If the user read protection key hash is programmed with a nonreset value, meaning that the read protection is enabled, or if the key

hash CRC is invalid, the SWD is disabled by the kernel and SWD access to the device is not possible. Perform this programming only after product development is complete and SWD access is not intended in the field. However, when read protection is enabled, the SWD is opened only if CRC protection is enabled and the CRC is corrupted, which allows device recovery when the CRC is accidentally corrupted. In this case, the SWD opens, but the user flash memory is protected with no read, write, or page erase accessibility (to maintain the user code confidentiality while allowing device recovery). However, mass erase is still possible, which results in the user flash memory being open again (with read and write access).

The user flash memory (user space) is completely blank when shipped. Therefore, none of the security keys and parameters are programmed. Most of the parameters, such as the key hash CRC and the user code length, have invalid values, meaning the parameters are all set to 0xFF. In this case, the kernel performs a check of the user flash memory to identify if it is blank or completely unprogrammed. If the user flash is blank, the kernel

skips all the checks and opens the SWD. In addition to opening the SWD, which allows users to connect to the device through SWD for their development, the boot kernel also enters the UART downloader mode and awaits reception of the SSL.

After the device is programmed via the SWD, the user flash memory is no longer blank, and the kernel relies on the state of the SYS_BMODE0 pin to decide if the user code must be executed after performing all the checks explained in the Boot Code Flow section, or if the boot kernel must enter the UART downloader mode.

If the SYS_BMODE0 pin is asserted low, the kernel enters the UART downloader and waits for the SSL to download.

If the SYS_BMODE0 pin is deasserted high, the kernel jumps to the user reset vector in the user flash memory after performing all security checks.

The only case where the kernel enters the UART download mode without sampling the SYS_BMODE0 pin is when the user flash memory is blank.




Figure 27. Boot Kernel Flowchart

SSL code downloaded over the UART must be mapped to the SRAM. In UART downloader mode, the SSL is loaded to the SRAM and has flash programming capabilities. The kernel authenticates the SSL and allows execution only if authentication is successful. This code is responsible for downloading and upgrading the actual firmware (for example, the user application) in the user flash memory. The kernel does not support direct updates to the user flash memory. Therefore, the SSL is required to perform such actions.

The kernel follows a specific protocol to download the SSL to the processor, which must be adhered to by the transmitting host. If the SSL follows the same packet protocol as the kernel, the host interface is simplified (for example, communication with the kernel and the SSL is uniform). The details of the protocol are discussed in the Protocol section.

UART DOWNLOADER

The ADuCM4050 processors enter UART downloader mode if the SYS_BMODE0 pin (GPIO17) is pulled low. If this condition is detected by the device at power-on or hard reset, the device enters serial download mode. In this mode, an on-chip loader routine in the kernel is initiated, which configures the UART port of the device and, via a specific serial download protocol, communicates with a host to manage the firmware upgrade process. Figure 28 shows the UART downloader flow.




Figure 28. UART Downloader Flowchart

16067-040

Protocol

After the serial downloader is triggered by asserting the SYS_BMODE0 pin, the kernel waits for the host to send a carriage return character (ASCII 0x0D, as shown in Figure 29) to initiate the UART autobaud process.

The kernel makes use of the UART autobaud feature to detect the baud rate of the host and to subsequently configure the UART port to transmit or receive, at the baud rate of the host, with eight data bits and no parity. Due to the 6.5 MHz reset peripheral clock (PCLK), the UART can be configured by the kernel to support baud rates up to 230,400 bps. Baud rates greater than this value contain more errors and may result in an unreliable data transfer. However, after loading the SSL, higher baud rates are possible if the SSL increases the PCLK (up to 52 MHz) and performs a second autobaud detection via the UART.

After receiving the autobaud character, the kernel calculates the required clock divisor values and configures the UART, at which point the kernel sends the device information as part of a 57-byte ID data packet, as shown in Figure 29, to acknowledge that the autobaud detection process is successful.

Packet Structure

In addition to indicating to the host that the processor is now ready to communicate, the autobaud acknowledgement also contains information about the device, the state of the user flash memory, and security restrictions. After the autobaud acknowledgement, the data transfer itself can begin, as governed by the communications data transport packet format shown in Table 16.




Figure 29. Autobaud Character

Table 6. Autobaud Response

Bytes	Contents
1 to 15	Product identifier: ADuCM4050 and six spaces
16 to 18	Hardware and firmware version numbers
19	User code blank; x means the code to execute, and a dash (-) means the user code is blank
20	User code checksum; P means that the checksum passed, and F means that the checksum failed
21	Write protection enabled; W means that write protection is disabled, and a dash (-) means that the write protection is enabled
22	Read protection enabled; R means that read protection is disabled, and a dash (-) means that read protection is enabled
23	Space
24 to 55	128-bit serial number, as a 32-digit uppercase hexadecimal number (for example, 0123456789ABCDEF0123456789ABCDEF)
56	Line feed
57	Carriage return

Table 7. UART Packet Structure

ID0	ID1	Number of Data Bytes	CMD	Value	Data	Checksum
0x07	0x0E	5 to 255	W, R, or I	h, u, m, l	xx	CS

Packet Start ID Field, ID0 and ID1

The first transfer field is the 2-byte packet start ID field (ID0 and ID1), comprised of two start characters (0x07 for ID0 and 0x0E for ID1). These bytes are constant and are used by the loader to detect the beginning of a valid data packet.

Number of Data Bytes Field

The next transfer field is the total number of data bytes field, which includes the 1-byte command (CMD), the 4-byte address (value), and the remaining payload (data). The minimum number of data bytes is five, which corresponds to a command and address only. The maximum number of data bytes is 255, supporting a command, an address, and up to 250 bytes of data.

Command Function Field (CMD), Data Byte 1

The command function field describes the function of the data packet. Three commands are supported by the kernel, represented in ASCII format:

- W (0x57)—write command
- R (0x52)—run command
- I (0x45)—information command

Write Command

The write command packet shown in Table 8 includes the number of data bytes ($5 + n$, where n is the payload size in bytes), the write command (W), the 32-bit start address to write to, and the n data bytes in the payload.

When a write command packet is received by the kernel, the payload bytes are placed sequentially in the SRAM as they arrive, beginning at the start address. The kernel sends a no acknowledge command if the checksum is incorrect or if the received address is out of range. If the host receives a no acknowledge from the loader, abort and restart the download process.

Run Command

After the host transmits all the data packets to the kernel, it can send a final packet instructing the kernel to start executing code.

Table 8. Write Command Packet

ID0	ID1	Number of Data Bytes	CMD	Value	Data	Checksum
0x07	0x0E	$5 + n$	W (0x57)	Start address	n bytes	CS

Table 9. Run Command Packet

ID0	ID1	Number of Data Bytes	CMD	Value	Checksum
0x07	0x0E	5	R (0x52)	Start address	CS

Table 10. Information Command Packet

ID0	ID1	Number of Data Bytes	CMD	Value	Checksum
0x07	0x0E	5	I (0x52)	0xFFFFFFFF	CS

This final packet is achieved by sending the run command packet, which is comprised of the run command (R) and the 32-bit address to begin running from, as shown in Table 9.

When the kernel receives a run command packet, it jumps to the address supplied in the packet only after the permission checks pass.

Information Command

The host can send the information command packet shown in Table 10 at any time. This packet is comprised of the command (I) and a 32-bit address. Though the value field is required for the packet to be properly received by the kernel, the content of the value field is irrelevant.

When the kernel receives the information command packet, it responds with the 57-byte ID packet (see Table 6).

Value Field (Data Byte 2 to Data Byte 5)

The value field contains a 32-bit address that includes h, u, m, and locations. The MSB is in the h location (Data Byte 2), and the LSB is in the l location (Data Byte 5).

The significance of the packets associated are as follows:

- In a write command packet, the value field indicates the start address in memory to which the data payload is written.
- In a run command packet, the value field indicates the address in SRAM where the SSL code begins.
- In an information command packet, the value field has no meaning.

Data Field (Data Byte 6 to Data Byte 255)

User code is downloaded one byte at a time, and the data field can contain a maximum of 250 bytes. The data is normally stripped out of the Intel® HEX extended 16-byte record format, reassembled by the host, and then sent in packet form using a series of write command packets to the ADuCM4050 processor.

Checksum Field (CS)

The data packet checksum is written to the checksum field. This twos complement checksum is calculated from the summation of the hexadecimal values that span the number of bytes field to the end of the data field. Thus, the 8-bit LSB of the sum of all the bytes in the packet from the number of data bytes field, up to and including the checksum field, is 0.

Acknowledge of Command

The loader routine issues a no acknowledge command (0x07) as a negative response, or an acknowledge command (0x06) as a positive response to each data packet received.

The loader transmits a no acknowledge if it meets any of following conditions:

- The loader receives an incorrect checksum.
- A UART framing or break error occurs (this error may not reach the host if the UART link is invalid).
- The SRAM code verification fails.

If any one of these conditions is met, it is required to reset the target and restart the firmware upgrade process. If none of these conditions are met, an acknowledge command is transmitted.

READ PROTECTION KEY AND HASHING

The read protection key allows access to the device during failure analysis. If the device is read protected and failure analysis of the current flash memory content is necessary, enable the SWD interface by sending the key corresponding to the hash stored in the user flash memory. It is recommended that the key be unique to the device and be based on the unique identifier of the device (for example, the serial number stored in the information space).

A hash is stored in the user flash memory after the interrupt vectors. This is the hash of a secret customer key. It is strongly recommended that this key be unique to the device for security

reasons, and that the unlock key is valid for that one specific device. To maintain a unique key per device, there must be a device identifier to associate which key belongs to a particular device. For simple key management, it is advised to make the key a hash of a master secret and the device identifier.

Read Protection Key = Hash (Master Secret || Unique Device Identifier)

Key Hash = Hash (Read Protection Key)

When the kernel is in UART loader mode, it can accept the read protection key. Then, the boot loader performs a hash of the read protection key and compares it to the stored key hash. On a successful match, the bootloader permits the downloaded SSL code in the SRAM to be executed with all the permissions enabled. If the key hash check fails, then the kernel checks the ICWP key in the user flash memory. If ICWP is turned off by the user by programming any value to Address 0x00000198 other than the hexadecimal equivalent of the NoWr ASCII string, then the SSL is allowed to run after protecting the flash against read and write accesses. In this case, the SSL must first issue a mass erase of the user flash memory before attempting to perform any access to the user flash memory space. If ICWP is also enabled by the user, then the SSL is not granted permission to run unless the key hash authentication passes.

The 128-bit read protection key is passed as a part of the SRAM code. This key must be stored in big endian format in the SRAM as a data payload starting at Address 0x20000180 and must be oriented in a specific fashion in the memory for the kernel to parse it correctly. Specifically, if the read protection key is represented as ABCDEFGHIJKLMNOP, where each letter represents one byte (with A being the first byte and P being the last byte), the required arrangement of the bytes in memory is shown in Table 11.

Table 11. Read Protection in SRAM

Address	Byte 0	Byte 1	Byte 2	Byte 3
0x20000180	D	C	B	A
0x20000184	H	G	F	E
0x20000188	L	K	J	I
0x2000018C	P	O	N	M

For example, if the read protection key is 0x00010203040506070-8090A0B0C0D0E0F, then Table 12 shows how the memory must be written.

Table 12. Example Read Protection Key in SRAM

Address	Byte 0	Byte 1	Byte 2	Byte 3
0x20000180	0x03	0x02	0x01	0x00
0x20000184	0x07	0x06	0x05	0x04
0x20000188	0x0B	0x0A	0x09	0x08
0x2000018C	0x0F	0x0E	0x0D	0x0C

The kernel computes the SHA-256 hash of this key, truncates it to a 128-bit hash, and then compares it to the hash stored in Page 0 of the user flash memory at Address 0x00000180. The user must store the 128-bit truncated hash of the key to the flash memory using a similar pattern. The SHA-256 hash for the example key shown in Table 12 is 0xBE45CB2605BF36BEBDE68-4841A28F0FD43C69850A3DCE5FEDBA69928EE3A8991, which means the 128-bit truncated hash that must be stored properly to the user flash memory space is 0x43C69850A3DCE-5FEDBA69928EE3A8991, arranged as shown in Table 13.

Table 13. Example Read Protection Key Hash in Flash Memory

Address	Byte 0	Byte 1	Byte 2	Byte 3
0x00000180	0x50	0x98	0xC6	0x00
0x00000184	0xFE	0xE5	0xDC	0xA3
0x00000188	0x28	0x99	0xA6	0xDB
0x0000018C	0x91	0x89	0x3A	0xEE

The CRC32 of the key hash is calculated with a polynomial of 0x04C11DB7 and a seed value of 0xFFFFFFFF, and it is stored in LSB format in the flash memory space at Address 0x00000190.

MEMORY CONFIGURATION

Table 14 summarizes the different keys and parameters stored in Page 0 of the user flash memory, the associated addresses, and the values programmed to Page 0 when creating a project with the default start-up file.

Table 14. Page 0 Memory Configuration

Content	Address Range		Size (Bytes)	Section Name	Default Content
	Start Address	End Address			
Vector Table	0x0000_0000	0x0000_017F	384	.intvec	Vector table
Read Protection Key Hash	0x0000_0180	0x0000_018F	16	ReadProtection KeyHash	0xFFFFFFFF 0xFFFFFFFF 0xFFFFFFFF 0xFFFFFFFF
CRC of Read Protection Key	0x0000_0190	0x0000_0193	4	CRC_ReadProtection KeyHash	0xA79C3203
Number of Pages the CRC Computes	0x0000_0194	0x0000_0197	4	NumCRCPages	0
Checksum	0x0000_07FC	0x0000_07FF	4	Checksum	Checksum of 0 to 0x7FB (if enabled in tools by the user)
Page 0 User Memory	0x0000_01A0	0x0000_07FC	1628	Page0_region	User application

HANDLING CRC IN THE IAR WORKBENCH

Calculate the CRC from part of the application image to be loaded into the first several pages of the flash memory. Store the page number of the last page involved in the CRC calculation at Address 0x194 as a 32-bit integer. For example, if only Page 0 is involved in the CRC calculation, store the value of 0x00 at Address 0x194. If the CRC is calculated for the first three pages, the value must be 0x02.

When the CRC is calculated, the last four bytes of the last page included in the CRC calculation are excluded. These four bytes are used for storing the CRC value itself. For example, if the last page is Page 0, the CRC is calculated from Address 0x000 up to and including Address 0x7FB. The tool stores the calculated CRC value at Address 0x7FC as a 32-bit integer.

The standard CRC calculation is CRC32 with a polynomial of 0x04C11DB7, stored in MSB first format, with an initial value of 0xFFFFFFFF. The unit size is 32 bits, which means the tool must read 32 bits at one time from the image when calculating the CRC.

Checksum Tab

There is a **Checksum** tab under the **Linker** category in the IAR tools, which generates the CRC of the user application code. To store the correct CRC, the following settings must be used (see Figure 30):

- Check the **Fill unused code memory** box.
- Set the **End address** field to 0x7FB (this value changes depending on the page number).
- Check the **Generate checksum** box.
- Select the **4 bytes** option from the **Checksum size** pull-down menu
- Set the **Alignment** field to 4 (which indicates 4 bytes).
- Select the **CRC32** option from the **Algorithm** pull-down menu.
- Set the **Initial value** field to 0xFFFFFFFF and ensure the **Use as input** box is not checked.
- Select the **32-bit** option from the **Checksum unit size** pull-down menu.




Figure 30. Checksum Settings

CROSSCORE SERIAL FLASH PROGRAMMER

The CrossCore® Serial Flash Programmer (CCSFP) is a PC-based host utility, provided by Analog Devices that upgrades the user code over the UART port. CCSFP provides a graphical user interface (GUI) to provide the following options for the UART upgrade:

- Target processor
- UART PC port number
- Baud rate
- SSL hexadecimal file to be used for the upgrade
- User application hexadecimal file to be upgraded
- Key to authenticate the SSL

Figure 31 shows the GUI for the CCSFP. The user must provide the SSL in the **Second stage kernel** field, which is first downloaded into the SRAM of the processor and then is executed before the user application in the **File to download** field is sent to flash, based on the authentication. The 128-bit key for the authentication can be entered in the **Key** field.




Figure 31. CrossCore Serial Flash Programmer GUI

The **Status** window shows the state of the UART download process, device related information, and the status of the commands as returned by the kernel. As shown in Figure 31, the **Status** window shows the device information sent by the kernel, showing the product ID, serial number, and user code status. After the SSL is downloaded, as indicated by the **Download completed** message displayed in the **Status** window, the SSL is then authenticated by the kernel and the actual user application is sent.

Figure 32 shows the SSL executing on the device, receiving the user application, and writing it to the user flash memory space.




Figure 32. User Application Code Being Written by the SSL to User Flash Memory via the CCSFP

CACHE MEMORY IN THE ADuCM4050

The memory in the ADuCM4050 has 512 kB of embedded flash memory within the error correction code (ECC), a 96 kB data SRAM with parity, and 32 kB user configurable instruction and data SRAM with parity. Four kB of SRAM can be used as cache memory to reduce active power consumption by reducing access to the flash memory.

The cache in the ADuCM4050 consists of a low power cache controller for the instruction code (ICODE) and data code (DCODE) accesses, a 4 kB instruction cache with two way associativity, and a line size of 256 bits. The instruction cache has a least recently used replacement policy. The cache writes to flash, and the core can issue writes to the flash only through the advanced peripheral bus (APB) interface of the flash controller. If the code is placed in the flash, enabling the cache helps the speed of execution. For more information, see the Effects of Cache on the Speed of Execution section. For details on current consumption, see the Current Consumption Comparison section.

This section discusses the use of an on-board cache controller to use a portion of the SRAM as instruction and data cache for user code that otherwise executes from the flash memory.

STORAGE AND ACCESS ECOSYSTEM BLOCK DIAGRAM

The ADuCM4050 cache architecture consists of a digital cache controller, a cache memory implemented as part of the system SRAM, a digital flash controller, and a flash memory. The cache architecture decreases the average latency of instruction and data accesses by utilizing the faster SRAM memory, and decreases the frequency of accesses to the relatively higher power flash memory.

When code is executed from the flash memory with the cache enabled, frequently used instructions are automatically cached in a dedicated region of the SRAM. In most applications, no further user effort is required, though locking and control features are provided.




Figure 33. Storage and Access Ecosystem Block Diagram

FLASH CONTROLLER

The flash controller is coupled with a cache controller module, which provides two advanced microcontroller bus architecture high performance bus (AMBA AHB) ports: one port for reading data (DCODE), the other for reading instructions (ICODE). The flash controller supports simultaneous ICODE and DCODE read accesses. DCODE has priority on contention.

The flash controller implements a prefetch mechanism to optimize ICODE read performance. This mechanism provides optimal performance when reading consecutive addresses on the ICODE interface. Simultaneous reads are possible if the ICODE read returns buffered data from prefetch.




Figure 34. Flash and Cache Controllers

EFFECTS OF CACHE

Effects of Cache on the Speed of Execution

The flash memory and the SRAM memory have distinct power and performance profiles.

The cache architecture copies a portion of user code into the SRAM during execution where instruction and data read latency is lower. For every instruction or data read that is satisfied by the cache memory, the overall system performance is improved.

Using the cache generally increases the speed of execution. The extent of increase depends on the type of code used. If the code has loops that fit completely into the cache, the speed of execution increases significantly, because the majority of instruction accesses is served from the faster SRAM memory. For a loop code that fits into the cache completely, the speed of access is 15% to 20% faster than using only flash. If the code is generally linear and/or jumps between segments too large to fit into the cache, the speed of execution is not significantly improved, because the majority of instruction accesses is served from the slower flash memory.

Each cache miss results in a cache line fill consisting of four 64-bit reads from the flash memory.

When the core is running at 52 MHz, the flash needs an extra two wait states to be enabled. As a result, flash accesses consume more cycles compared to that at 26 MHz. Thus, cache may have more of an advantage in this case. A comparison table of the cycle count consumption for a prime number code (loop code that can fit in the cache) is as shown in Table 15.

Table 15. Cycle Count Comparison

Code Executing	Cycle Bount at 26 MHz	Cycle Count at 52 MHz
Flash Only	68,875	107,676
Flash + Cache	57,137	57,137

As shown in Table 15, there is about a 17% decrease in the cycle consumption in flash and cache compared to flash only at 26 MHz and about a 47% decrease in the cycle consumption in flash and cache compared to flash only at 52 MHz, due to the extra two wait states of flash.

Cache Key Register

Address: 0x40018060, Reset: 0x00000000, Name: FLCC0_CACHE_KEY




Table 16. Bit Descriptions for FLCC0_CACHE_KEY

Bits	Bit Name	Description	Reset	Access
[31:0]	VALUE	Cache Key Register. Enter 0xF123_F456 to set the user key. Returns 0x0 if read. The key is cleared automatically after writing to FLCC_SETUP register.	0x0	W

Cache Setup Register

Address: 0x04001805C, Reset: 0x00000000, Name: FLCC0_CACHE_SETUP

The cache user key is required to enable a write to this location. The key is cleared after a write to this register.




Table 17. Bit Descriptions for FLCC0_CACHE_SETUP

Bits	Bit Name	Description	Reset	Access
[31:1]	RESERVED	Reserved.	0x0	R
0	ICEN	I-Cache Enable. If this bit is set, then I-Cache is enabled for AHB accesses. If 0, then I-Cache is disabled, and all AHB is accessed via flash memory.	0x0	R/W

The details of the FLCC0_CACHE_SETUP register are shown in this section. Follow these instructions to enable or disable the cache:

1. The instruction cache (ICACHE) is disabled by default. To enable the ICACHE or toggle, the 0xF123F456 key must be written into the FLCC0_CACHE_KEY register.
*pREG_FLCC0_CACHE_KEY = 0xF123F456;
2. To enable the ICACHE, set the ICEN bit in the FLCC0_CACHE_SETUP register. Clear this bit to disable the ICACHE.

*pREG_FLCC0_CACHE_SETUP |= (1 << BITP_FLCC_CACHE_SETUP_ICEN);

Effects of Cache on Current

When it comes to current consumption, SRAM accesses consume less current than flash accesses. Therefore, when the cache is enabled during the execution of code from the flash memory, the current consumption is generally between that of code executing directly from the flash or directly from the SRAM, unless the code is such that every cache access is a miss. In this case, the current is higher than executing from flash alone. Cache line fills result in approximately 2× more flash reads than occur if executing directly from the flash. This read rate is 4× more, if not for the prefetch buffer in the flash also performing a read, which must miss too, if the cache misses.

$$Current_{SRAM} \leq Current_{CACHE}$$

When using the cache, the current consumption is proportional to the cache miss rate. This result is because of a scalar current reduction for each cache hit, because the data or code access is served from the lower current SRAM, rather than the higher current flash memory. Therefore, code consisting of many small loops is more greatly affected than linear code or code consisting of segments too large to fit into the cache memory.

The cache usage can also be inferred from the current consumption. If the current consumption using cache and flash is nearer to the current consumption when using the SRAM, the

cache hit rate must be high. If the current consumption using cache and flash is nearer to the current consumption when using only the flash, the cache hit rate must be low.

CURRENT CONSUMPTION COMPARISON

For a loop code (prime number code in this example), it is seen that the current consumption using the flash and cache (1.47 mA) is close to the current consumption in the SRAM (1.30 mA). This result is because the code consists of many loops, each small enough to fit into the cache. Therefore, the accesses to the flash are relatively rare, and the code mostly executes from the relatively low power cache. Accesses are minimal and, consequently, a minimal increase in the current consumption in flash and cache compared to the current consumption using the SRAM is seen.

For a linear code (ULPBench code in this example), it is seen that the current consumption using the flash and cache deviates away from the current consumption seen using only the SRAM. This is because the ULPBench code is mostly linear and does not fit well into the cache. Therefore, there is still a similar number of accesses into the flash memory as there are when executing directly from the flash. The cache misses are too high. As a result, the current consumption deviates away to a greater extent. The measurements shown in Table 18 are with HCLK at 26 MHz, peripheral clock off an HP buck switched on.

Table 18. Current Consumption Comparison

Type of Code	SRAM (mA)	Flash (mA)	Flash and Cache (mA)	Cache Misses (~12 sec of Execution)
Loop Code (Prime Number)	1.30	1.85	1.47	24
Linear Code (ULPBench)	2.83	3.34	3.12	~800,000

DUAL RTC FEATURE IN THE ADuCM4050

In many applications, an RTC time stamps sensor data. The RTC must run even when the MCU is in a deep sleep mode. A low power RTC is crucial to achieving a long battery life.

Notable features of ADuCM4050 RTCs include the following:

- A dual RTC (RTC0 and RTC1). Both RTCs can be used as wake-up timers.
- SensorStrobe and input capture features.

This section provides guidelines for choosing between RTC0 and RTC1, depending on the power modes and functionality required.

COMPARISON OF THE RTC FEATURES

The ADuCM4050 has two RTCs, RTC0 and RTC1 (also named FLEX_RTC). Table 19 shows differences between both RTCs.

POWER CONSIDERATIONS

Table 20 shows current consumption when using RTC0 and RTC1 in different use cases. Four scenarios are considered as follows:

- Scenario 1. The device switches between the active and hibernate power modes, and the application requires high time accuracy. Either RTC0 or RTC1 can be used in this scenario, but using RTC1 is recommended from a power point of view, because RTC1 uses less power.

- Scenario 2. The device switches between the active and hibernate power modes, and the application does not require high time accuracy. Either RTC0 or RTC1 can be used in this scenario, but using RTC1 is recommended from a power point of view, because RTC1 uses less power.
- Scenario 3. The device switches between the active and shutdown power modes. Only RTC0 can be used in this scenario, because RTC1 is not active in shutdown mode.
- Scenario 4. The device switches between the active, shutdown, and hibernate power modes. Only RTC0 can be used, in this scenario, because RTC1 is not active in shutdown mode.

A basic program comprised of an RTC alert to wake up the ADuCM4050 microcontroller from low power mode and to toggle an LED measured the deep sleep power modes current.

CONCLUSION

Use RTC0 in applications that use shutdown mode and require an RTC.

RTC1 is a feature rich RTC that enables ultra low power consumption in applications that do not use shutdown mode. Typical applications for which RTC1 is suited are applications in which the ADuCM4050 microcontroller sends output pulses to external sensors via a general-purpose input/output. The SensorStrobe mechanism is only available in RTC1.

Table 19. Summary of the Differences Between RTC0 and RTC1

Feature	RTC0	RTC1
Resolution of the Time Base (Prescaling)	RTC0 counts time at 1 Hz in units of seconds only.	RTC1 can prescale the clock by any power of 2 from 1 to 15, counting time in units of any of these 15 possible prescale settings.
Wake-Up Timer	The wake-up time is specified in units of seconds.	The wake-up time can be specified in units of any power of 2 multiple of 30.7 μ s up to 1 second.
Number of Alarms	One alarm only, which uses an absolute, nonrepeating alarm time.	Two alarms: one absolute alarm time and one periodic alarm, repeating every 60 prescaled time units.
Power Domain	Powered off V_{BAT} domain and is always on; RTC0 can function in all power modes.	Powered off 1.2 V (VREG) domain; RTC1 can function in all power modes except shutdown mode.
SensorStrobe and Input Capture Features	Not supported.	Supports four input capture channels and four SensorStrobe channels.
Source Clock	Low frequency crystal (LFXTAL).	Depending on the low frequency multiplexer (LFMUX) configuration, the RTC is clocked by LFXTAL or the low frequency oscillator (LFOSC).

Table 20. Comparison of Current Consumption in Different Use Case Scenarios

Scenario Number	Use Case ¹	Recommended RTC	Sleep Mode Current (nA)
1	Active to hibernate	RTC1 (LFXTAL)	783
2	Active to hibernate	RTC1 (LFOSC)	720
3	Active to shutdown	RTC0 (LFXTAL)	387
4	Active to hibernate	RTC0 (LFXTAL)	783

¹ The device switches between the modes listed in this column.

BENEFITS OF ADuCM4050 THE DC-TO-DC CONVERTER

This section discusses the advantages and disadvantages of charge pump converters vs. inductor converters, the latter of which are frequently used. This section demonstrates why this architecture is used on the ADuCM4050 microcontroller, accounting for advantages in many aspects including price, area, simplicity, and ease of use.

Direct current-to-direct current (dc-to-dc) converters are key blocks in designs where it is required to manage different voltage domains, such as in the ADuCM4050 microcontroller.

Methods of dc-to-dc conversion are briefly explained in this section to provide users with context. A charge pump converter is chosen for use in the ADuCM4050 because of its advantages when compared to other configurations.

The purpose of this section is to help users understand why the capacitive dc-to-dc converter is a better alternative to inductive conversion solutions in the ultra low power applications for which the ADuCM4050 is intended.

This section provides details and examples to prove the qualities and benefits of this charge pump converter solution. Figure 35 shows the buck enabled design present in the ADuCM4050 microcontrollers. The ADuCM4050 uses a charge pump converter, which is not used in the majority of microcontrollers with similar characteristics available on the market. Other microcontrollers usually use traditional inductor converter architectures.




Figure 35. ADuCM4050 Buck Enabled Design

DC-TO-DC BASICS

The ADuCM4050 processor is intended for ultra low power applications. Power efficiency is one of the key considerations in such applications. Therefore, using a dc-to-dc converter is crucial in designs in which power must be used as efficiently as possible.

There are different ways to perform dc-to-dc voltage conversions. Such conversions involve stepping up or stepping down the dc voltage that power the device.

DC-to-DC Conversion Methods

The most extensively used methods for regulating the different power domains of a system are switching conversion and linear regulation. Select the method that best meets the requirements of the design or application.

Linear Regulators

Linear regulators consist of a network of resistive dividers that dissipate excess voltage. Linear regulators are widely used due to the ease of use and implementation, as well as the low cost.

In ultra low power applications, linear regulators are less efficient when compared to switching converters. In a linear regulator, the output current is approximately the same as the input current, and its operating principle is to dissipate any leftover voltage. Switching converters perform the same action more efficiently.

The ADP165/ADP166 devices are very low quiescent current, LDO, linear regulators. The ground current represents the difference between input and output currents. Figure 36 represents the ADP165/ADP166 ground current vs. the load current (I_{LOAD}), showing the small difference in currents in a linear regulator.




Figure 36. ADP165/ADP166 Ground Current vs. Load Current (I_{LOAD})

To analyze the power efficiency of the solution, consider a typical application based on linear regulators. For an input voltage of 3 V, an output voltage of 1 V, and an output current of 1 μ A, the input current is approximately 1 μ A. This scenario results in 33% efficiency (see Equation 3).

$$\text{Efficiency} = \frac{\text{Energy Output}}{\text{Energy Input}} \times 100\% \quad (1)$$

$$\text{Efficiency} = \frac{I_{OUT} \times V_{OUT} \times t}{I_{IN} \times V_{IN} \times t} \times 100\% \quad (2)$$

where:

I_{OUT} is output current.

V_{OUT} is output voltage.

t is time.

I_{IN} is input current.

V_{IN} is input voltage.

$$\text{Linear Regulator Efficiency} = \frac{1 \times 1 \times t}{1 \times 3 \times t} \times 100\% = 33\% \quad (3)$$

Using a switching converter instead of a linear regulator, the input current is 1/3 μ A, leading to 100% efficiency in an ideal performance, as shown in Equation 4.

$$\text{Efficiency} = \frac{1 \times 1 \times t}{1/3 \times 3 \times t} \times 100\% = 100\% \quad (4)$$

In general, switching converters are more efficient than linear regulators. Moreover, losses in efficiency produce an increase in temperature that is much higher in linear regulators because their dissipation must be larger to achieve the same conversion. Additionally, linear regulators require more investment in management to reduce temperature.

Traditionally, Analog Devices uses linear regulators in designs that precede the ADuCM4050 microcontrollers because of their simplicity and low cost. Currently, it is common to locate linear regulators at the output of charge pump converters to stabilize their rippled output.

Switching Converters

Switching converters use switches and components with low losses, such as inductors or capacitors, to regulate voltage. Typically, these components are charged and discharged by switching transistors. This section discusses two types of switching converters: charge pump converters and inductor converters.

Inductor converters are among the most commonly used converters in microcontroller designs to achieve ultra low power with high efficiency. This efficiency and the wide gain range make this architecture desirable.

The charge pump, or switched capacitor converter, is an alternative to inductive converters. The charge pump process is performed by connecting and disconnecting switches to charge and discharge capacitors. This process is achieved without inductors, which saves space and costs.

The ADP2503/ADP2504 are high efficiency inductor converters that can operate at input voltages greater than, less than, or equal to the regulated output voltage. The ADM660/ADM8660 are charge pump voltage converters that can achieve efficiency greater than 90% with low output currents (up to 50 mA). Figure 37 and Figure 38 show the efficiency for a given input voltage and output currents for the ADP2503/ADP2504 devices and ADM660/ADM8660 devices, respectively.

As observed in these Figure 37 and Figure 38, charge pump converters are less efficient than inductor converters due to the output shape for output load currents. In contrast, charge pump converters are an appropriate solution to low load current applications.




Figure 37. ADP2503/ADP2504 Efficiency vs. Output Current (I_{OUT})




Figure 38. ADM660/ADM8660 Efficiency vs. Load Current

The ADuCM4050 microcontrollers have a linear regulator at the charge pump converter output to adjust and stabilize the supply of the digital core and memories. Furthermore, the devices have the ability to bypass the charge pump converter to only use the linear regulator to reduce and adjust the voltage. This feature allows the user to decide between using a traditional solution or to improve efficiency and increase power savings using the charge pump block at the expense of two extra 0.1 μ F capacitors.

In general, inductor converters do not require a linear regulator at their output, which is inconvenient in charge pump converter designs. However, some microcontrollers available on the market that use an inductor converter solution include a linear regulator in the inductor converter. Despite this disadvantage, charge pump converters offer a breadth of advantages to be considered, as follows:

- Area
- Thickness
- Price
- Design simplicity
- Ease of use
- Electromagnetic interference (EMI)

The Capacitors vs. Inductor Converters section discusses each advantage of charge pump converters when compared to inductor converters.

CAPACITORS VS. INDUCTOR CONVERTERS

Area

Charge pump converters do not require inductors to accomplish dc-to-dc voltage conversion, whereas inductor converters require inductors, capacitors, and other components, such as resistors, to fulfill this task. This fact allows the design of smaller printed circuit boards (PCBs), saving area and cost.

This section compares the components required for using the ADuCM4050 charge pump converter against inductor converters used in other similar solutions available on the ultra low power microcontroller market. The components required in these architectures are included in their respective data sheets or user guides. In case of microcontrollers with on-chip inductor based converter(s), 0806 package inductor(s) are typically recommended for maximum power efficiency.

The two bill of materials shown in Table 21 and Table 22 demonstrate that charge pump converter covers less area than the inductive converters. The dimensions of a 0402 (1005 metric) package capacitor (length × width × thickness) are 1 mm × 0.5 mm × 0.55 mm. The dimensions for an 0806 (2016 metric) package inductor are 2 mm × 1.6 mm × 1 mm.

Table 21. Area of Charge Pump Converter Components in the ADuCM4050

Component	Value	Package	Quantity	Area (mm ²)
Capacitor	0.1 μ F	0402	2	1
Capacitor	1 μ F	0402	1	0.5
Total				1.5

Table 22. Area of Inductor Converters Components

Component	Value	Package	Quantity	Area (mm ²)
Capacitor	1 μ F	0402	2	1
Inductor	2.2 μ H	0806	2	6.4
Total				7.4

The area in a charge pump converter is smaller, because this type of converter employs fewer and smaller components than inductive solutions. Inductor converters use nearly five times the area of charge pump converters in terms of external components used.

In inductor converters solutions, the thickness of the circuit board is often determined by the inductors, because inductors are thicker than capacitors—1 mm vs. 0.55 mm.

Price

Issues to consider when working with inductor converters include the large number of components and the cost of these components.

Table 23 and Table 24 show the Bill of Materials (BOM) cost of the ADuCM4050 solution and an inductor based solution, respectively. Prices listed are indicative and at order quantities of 1 component.

Table 23. Price of Charge Pump Converter Components for the ADuCM4050

Component	Value (μ F)	Quantity	Price (USD)
Capacitor	0.1	2	0.26
Capacitor	1	1	0.21
Total			0.47

Table 24. Price of Inductor Converters Components

Component	Value	Quantity	Price (USD)
Capacitor	1 μ F	2	0.42
Inductor	2.2 μ H	2	0.70
Total			1.12

When considering both lists of materials, the difference in the total BOM is around 70 cents (USD). This amount, though seemingly small, produces a notable impact when multiplying the cost across many products; the difference in price is incremented because, in inductor designs, more components are present, and the engaged area is wider. Allowing smaller PCB designs reduces cost.

Compromising on inductor quality is possible to reduce cost, but doing so may lead to increased power dissipation and a degradation in efficiency.

There are clear advantages in both area and BOM cost when comparing charge pump converters and conventional inductor-based solutions.

Efficiency

As with the opportunity to reduce area, there also exists an opportunity to improve integration employing embedded components. This is a suitable scenario for charge pump converters to enhance efficiency, rather than inductors.

It is thought that charge pump converters are less efficient than inductor converters, which can be true when input voltages and loads change.

In charge pump converters, load changes are not a problem in ultra low power applications where low loads are managed.

Optimal efficiency is achieved with low load currents. Charge pump converters perform proper efficiencies with low loads, which is easily achieved by applying integration in an ultra low power application. The lower the required load, the better integration and efficiency are in charge pump converters.

By setting the appropriate configuration, charge pump converters are able to change their gain according to the input/output voltage ratio (V_{IN}/V_{OUT}). This process improves efficiency to achieve the same performance available with inductor converters.

Inductor based solutions use pulse-width modulation (PWM) to adjust the duty cycle to achieve suitable gain. Through this regulation, high efficiency is obtained, which decreases when the load lowers. Noise effects also appear during PWM, which results in increased cost for more expensive inductors.

If integration is required in inductor converter designs, embedded inductors require high frequency switching to work. High frequency switching results in power dissipation and efficiency losses, which is an undesirable outcome.

Electromagnetic Interference (EMI)

In charge pump converters, electromagnetic emissions are not relevant. Such radiation is not a cause for concern, unlike inductor magnetic radiation.

EMI is inconvenient when using inductors, even more so if they are switched inductors with behavior similar to an emitting antenna. Unpredictable interferences can occur in other parts of

the design or the evaluation board. Furthermore, it becomes a sensitive problem if radio frequency tasks are being performed.

Inductor converters replace PWM with pulse frequency modulation (PFM) when low loads are required to improve efficiency. If PWM is performed, switching noise and output voltage ripple are easily improved by a simple filter at the output voltage of the converter. However, the PFM method has a variable frequency band and may produce the resonance frequency of the filter. This wide frequency spectrum also results in high EMI.

CONCLUSIONS

Inductor converters are not suitable solutions in many senses when considering ultra low power applications. Inductor converters lose efficiency as load decreases. Their area is larger, which can lead to expensive components and greater costs. Inductor radiation poses a problem because EMI is more likely to occur at low loads, and so on.

Table 25 summarizes the advantages and disadvantages of three types of dc-to-dc converters. Evaluate the specific type of converter that best suits the application in question.

In conclusion, charge pump converters are the best solution in ultra low power applications. While other solutions worsen with low loads, charge pump converters are even better than in other situations.

Table 25. Comparison of Different Types of DC-to-DC Converters

Type of Converter	Advantages	Disadvantages
LDO	Simple Low cost No inductor No EMI	Less efficient than charge pumps and inductives
Charge Pump	Simple Low cost No inductor Cheaper than inductive Low loads Small area More efficient than LDO Low EMI	Less efficient than inductive at high loads EMI (less than inductive)
Inductive	Most efficient (not in low loads)	EMI Area Cost Poor efficiency at low loads Complex design

UART SOFTWARE FLOW CONTROL

Flow control is the process of managing the rate of data transmission between two nodes to prevent a fast transmitter from overwhelming a slow receiver. Flow control provides a mechanism for the receiver to control the transmission speed, so that the receiving node is not overwhelmed with data from the transmitting node.

UART flow control is a method for slow and fast devices to communicate with each other over the UART without the risk of losing data. Consider the case in which two units communicate over the UART. A Tx sends a long stream of bytes to a Rx. Rx is a slower device than Tx, and at some point, Rx cannot keep up with the speed of the data being transmitted. Therefore, Rx must either process some of the data or empty buffers before it can continue to receive data. Rx must instruct Tx to stop transmitting until Rx is ready to accept data. This method of waiting to transmit is known as flow control.

Flow control requires extra signaling to inform the transmitter to stop (pause) or start (resume) the transmission. The traditional hardware flow control in UART requires two extra signals: request to send (RTS) and clear to send (CTS). The logic level on these signals defines whether the transmitter continues to send data or must stop sending data. With software flow control, special characters are sent over the normal data lines to start or stop the transmission, thus using fewer signals.

This section describes the UART software flow control mechanism using the [ADuCM4050](#).

UART FLOW CONTROL

Hardware Flow Control

The hardware flow control mechanism uses out of band signaling to control the flow of data. In addition to the data signals, two extra signals—RTS and CTS—are required. These flow control signals are cross coupled between the two devices, with RTS on

one device being connected to CTS on the remote device, and vice versa as shown in Figure 40.

Each device uses the RTS to signal if it is ready to accept new data and reads the CTS signal to check if it is allowed to send data to the other device. As long as a device is ready to accept more data, the RTS signal is kept asserted. The device deasserts the RTS signal when its receive buffer is full.

The other device is required to respect the flow control signal and pause the transmission until the RTS signal is asserted again.

The flow control is bidirectional, meaning that both devices can request a halt in transmission. If one of the devices never requests a stop in transmission (for example, if the device is fast enough to always receive data), the CTS signal on the other device can be tied to the asserted logic level. Thus, the RTS pin on the fast device can become free to perform other functions.

Software Flow Control Using the XON and XOFF Signals

Software flow control does not require extra out of band signals. Only three signals are required: Rx, Tx and ground as shown in Figure 39. Software flow control is achieved by using special control flow characters. The control flow characters are sent over the normal Tx and Rx lines. These characters are typically ASCII codes, specifically XON (0x11) and XOFF (0x13), for resuming and halting the transfer, respectively.

If Device A sends XOFF to Device B, Device B halts transmission to Device A until Device B receives an XON character from Device A. If the data contains the XON and/or XOFF character, insert an escape character before the XON and/or XOFF character. The escape character used in this case is \ with ASCII Value 92 (0x5C). When this escape character is encountered, the character following it is considered to be a data character, not a flow control signal. If the data itself contains an escape character, ensure another escape character precedes the present escape character that is present in the data.




Figure 39. Software Flow Control Block Diagram




Figure 40. Hardware Flow Control

Sequence Diagram

Consider data communication between two devices—the ADuCM4050 MCU and a peer—where the MCU is transmitting and the peer is receiving. If the peer is slower than the MCU, the data transmission overwhelms the peer. At this stage, the peer sends an XOFF character to pause the transmission until the peer is able to process the data again. The MCU waits to receive a XON character from the peer. When the peer is ready to receive the data, it sends an XON character, instructing the MCU to resume transmission. In this way, using software flow control ensures that no data is lost. Figure 41 shows the sequence diagram of this described communication.




Figure 41. Software Flow Control Sequence Diagram

SYSTEM DESCRIPTION

Demonstration of UART software flow control using the ADuCM4050 is performed using the EV-COG-AD4050 evaluation kit. A PC with a terminal program running (such as HyperTerminal) is connected to the EV-COG-AD4050 UART port.




Figure 42. Connection Diagram

Handling Flow Control Signals from a Peer Device

The ADuCM4050 BSP contains drivers for all the peripherals, including UART. The software flow control mechanism is implemented in addition to the available UART driver functions.

The `adi_uart_Write_fc` function sends the XOFF and XON characters, and the UART interrupt service routine processes the XON and XOFF signals received from the PC.

adi_uart_Write_fc Function

When a write is issued using the `adi_uart_Write_fc` function, the global `RECV_XON` flag is checked to be aware whether the peer is ready to accept data. If the `RECV_XON` flag is false, it means that an XOFF signal is received, the peer cannot accept data, and, therefore, a failure is returned. If the `RECV_XON` flag is true, the peer is ready to accept data. The data is transmitted and a success is returned.

Example Code for Flow Control

The following code transmits data using flow control:

```

ADI_UART_RESULT adi_uart_Write_fc(
ADI_UART_HANDLE const hDevice, void *const
pBuffer, uint32_t nBufSize)
{
    /* Return code */
    ADI_UART_RESULT eResult;
    /* If there is no XOFF received, safe to
    transmit data */
    if(RECV_XON == true)
        eResult = adi_uart_Write (hDevice,
pBuffer, nBufSize);
    /* If XOFF is received, return fail */
    else
        eResult = ADI_UART_FAILED;
    return eResult;
}
  
```

Figure 43 shows the design of `adi_uart_Write_fc` function. When a write is issued, the `RECV_XON` flag is checked and, if the flag is true, the write is processed. If the flag is not true, it returns a failure. The `RECV_XON == TRUE` block in Figure 43 indicates the checking of the `RECV_XON` flag.




Figure 43. Flowchart of the `adi_uart_Write_fc` Function

Processing Control Signals from the Peer Through an Interrupt Service Routine (ISR)

The data received through the UART is monitored to check if it is a control signal or data. If the data received is an escape character, an escape flag (`bEscFlag`) is asserted so that the data following it is to be considered data and not as a control signal. If an XOFF or XON signal is received, it is checked to confirm if the escape flag is set. If the escape flag is not set, a global flag (`RECV_XON`) is updated.

When receiving an XOFF signal without the escape flag set, the `RECV_XON` flag is deasserted, meaning that it received an XOFF signal and data transmission must not happen. In the same way, the `RECV_XON` flag is asserted when receiving an XON signal without the escape flag set.

Data Processing Code Example

The following code processes the data that is received:

```

switch (readVal)
{
    /* If an escape is received */
    case FCEscape:
        /* If escape already received,
           consider it as data */
        if(bEscFlag == true)
            bEscFlag = false;
        else
            bEscFlag = true;
        break;
    /* If an XON is received */
    case XON:
        /* If escape received before,
           consider it as data */
        if(bEscFlag == true)
            bEscFlag = false;
        /* Valid control signal,
           update send flag */
        else
            RECV_XON = true;
        break;
    /* If an XOFF is received */
    case XOFF:
        /* If escape received before,
           consider it as data */
        if(bEscFlag == true)
            bEscFlag = false;
        /* Valid control signal,
           update send flag */
        else
            RECV_XON = false;
        break;
    default:
        break;
}
  
```

Figure 44 shows the design of the algorithm that handles and processes the flow control signals from the peer. The data received is first checked for an escape character. If an escape character is found, the data following it is considered to be data and not a

control signal. The received data is then checked for XON and XOFF control signals and the global RECV_XON flag is updated accordingly. The gray blocks in Figure 44 indicate the updating of the RECV_XON flag.




Figure 44. Flowchart of the Control Signals in the ISR Being Processed

16007-098

DATA CAPTURE

In the setup for data capture, the MCU is connected to a PC and communicates with a terminal program running on the PC. A UART sniffer, such as the serial port monitor, monitors the data communication occurring at a Baud rate of 9600. The data capture is performed using the sniffer, as discussed in this section.

Handling Flow Control Characters in the ADuCM4050

Figure 45 shows an example of handling the flow control signals from the peer.

Time	Direct...	Data	Data (chars)
00:575	UP	41	A
00:000	DOWN		
01:919	UP	42	B
00:000	DOWN		
01:934	UP	43	C
00:000	DOWN		
01:934	UP	44	D
00:000	DOWN		
01:918	UP	45	E
00:000	DOWN		
01:935	UP	46	F
00:000	DOWN		
00:514	DOWN		
00:000	UP	13	.
01:30:829	DOWN		
00:000	UP	11	.
00:028	UP	47	G
00:000	DOWN		
00:336	UP	41	A
00:000	DOWN		
01:757	DOWN		
00:000	UP	5c	\
00:178	UP	42	B
00:000	DOWN		
00:550	DOWN		
00:000	UP	13	.
01:369	UP	43	C
00:000	DOWN		
01:303	DOWN		
00:001	UP	5c	\
00:632	UP	44	D
00:000	DOWN		
00:087	DOWN		
00:000	UP	11	.
01:848	UP	45	E
00:000	DOWN		

Figure 45. Data Capture Using a Sniffer Program

Controlling the Received Data Flow

A simple procedure is implemented when controlling the received data flow to send the control signals from the ADuCM4050 MCU. In this case, the MCU is slower compared to the peer. The mechanism to send the control signals from the MCU is application specific, and the user can write an algorithm for sending the control signals.

As shown in Figure 46, an XOFF signal is sent after every five transmissions sent from the MCU. An XON signal is sent after a short interval of time. This implementation is an example, and it is described only for demonstration purposes. The user can develop a mechanism to handle the data and to send XON and XOFF signals.

Time	Direct...	Data	Data (chars)
17:098	UP	41	A
00:000	DOWN		
01:726	UP	42	B
00:000	DOWN		
01:711	UP	43	C
00:000	DOWN		
01:711	UP	44	D
00:000	DOWN		
01:727	UP	45	E
00:000	DOWN		
01:712	UP	13	.
00:000	DOWN		
17:148	UP	11	.
00:000	DOWN		
01:711	UP	46	F
00:000	DOWN		
01:728	UP	47	G
00:000	DOWN		
01:711	UP	48	H
00:000	DOWN		
01:712	UP	49	I
00:000	DOWN		
01:711	UP	4a	J
00:000	DOWN		

Figure 46. ADuCM4050 Transmitting Control Signals

SPI FLOW CONTROL METHODS

The SPI is an industry standard, synchronous serial link that allows full duplex operation to other SPI-compatible devices.

The ADuCM4050 SPI has enhanced modes of operation that provide the user the flexibility of half duplex operation and flow control options. The SPI data transfers use DMA transactions, allowing the ADuCM4050 core to be in sleep mode. Along with multibyte transfers in half duplex mode, this reduction in power consumption offers power savings that are essential for battery-powered designs, such as in wireless sensor networks.

Some notable features of the ADuCM4050 SPI are as follows:

- Continuous transfer mode.
- Read command mode for half duplex operation.
- Flow control.
- $\overline{CS}$ software override.
- Support for 3-pin SPI master or slave mode.
- LSB first transfer option.
- Interrupt mode. An interrupt is available after 1, 2, 3, 4, 5, 6, 7, or 8 bytes.

This section provides an understanding of the read command mode and flow control methods. These methods help lower the system power consumption when used with SPI slaves such as sensors, serial flash devices, ADCs, and RF transceivers.

SPI READ COMMAND MODE

Standard SPI masters communicate with slaves using the serial clock (SCK), master out, slave in (MOSI), master in, slave out (MISO), and chip select ($\overline{CS}$) lines as shown in Figure 47. The SCK, MOSI, and MISO signals can be shared by slaves, whereas each slave has a unique $\overline{CS}$ line. During an SPI transfer, data is simultaneously transmitted and received. The serial clock line synchronizes shifting and sampling of the information on the two serial data lines.

SPI transfers are typically full duplex. The transfers are controlled by the master. To receive data from the slave, the master must provide the clock, which is typically initiated when the data must be sent on the MOSI line.

Most SPI slaves mandate a protocol that must be used by the master for successful communication. The protocol can be as simple as a command, followed by an address (optional) and data (optional).

For example, a write command is unidirectional and typically involves the master transmitting the command, address (optional), and the data to be written to the address in the slave.

A read command requires the master to transmit the command and address (optional) and then reads the data associated with the address from the slave. If the data is multibyte, then the software on the master must write dummy data on the MOSI, which keeps the clock alive, to successfully read all the data bytes.

However, some SPI slaves require that, after the transmission of the read command byte on the MOSI, the data be read on MISO in a single $\overline{CS}$ transaction. An example of this requirement is shown in Figure 49.

The ADuCM4050 provides the read command mode to support such half duplex operations. The read command mode helps reduce the burden on the software and thereby the core execution cycles. In this mode, the user must specify the number of bytes to be transmitted and the number of bytes to be received in a transaction. It is also possible to specify if the data on the MISO must be ignored when the transmission on MOSI is in progress.

Using the read command mode allows the user to transmit a single byte and receive a set of data bytes from the slave, which is useful when the slave is a sensor or ADC providing a set of measured and processed data.

An application scenario is described in the System Description section, wherein the ADuCM4050 is the SPI master and a serial flash, W25Q32, is the SPI slave, as shown in Figure 48. Read command mode is helpful when pages of data must be read from the flash memory.




Figure 47. SPI Signals




Figure 48. Application Block Diagram




Figure 49. Read Command Mode

System Description

To showcase the read command mode, the following setup is used:

- Firmware—power-on self test application from the [ADuCM4050 BSP](#) for IAR.
- Hardware—[EV-COG-AD4050](#) board.

An oscilloscope is connected to the SPI lines to capture the signals. The oscilloscope plots, Figure 50 to Figure 54, show the SPI transfer between the [ADuCM4050](#) as the SPI master, and a serial flash W25Q32 as the SPI slave.

It is up to the user application to decide the transactions in which to use the read command mode.

Without Read Command Mode

In the reference application, the erase process of a 4 kB sector of the flash memory does not use read command mode. The absence of this mode can be observed from the transfer of the erase command and the address in individual chip select frames in Figure 50 and Figure 51.





Figure 50. Sector Erase

Figure 51 shows a single $\overline{CS}$ frame capture in the entire erase sequence. The $\overline{CS}$ line is toggled for the transfer of every command byte.

Figure 51. Sector Erase—Single $\overline{CS}$ Frame Capture

With Read Command Mode

Figure 52 shows one page read from the external flash. The size of one page of the flash is 256 bytes. This read sequence uses the read command mode and the entire read of the page happens in one $\overline{CS}$ transaction.




Figure 52. Page Read Sequence

Figure 53 shows the start of the page read sequence where the [ADuCM4050](#) transfers the command and address bytes. This transfer is followed by the page data from the serial flash memory.




Figure 53. Page Read Start Sequence (Command and Address Bytes)

Figure 54 shows a single byte read, which is part of the page read sequence.




Figure 54. Page Read—Single Data Byte

FLOW CONTROL MODES

Flow control is necessary to synchronize the data flow between a master and slave. The ADuCM4050 provides flow control as a differentiating feature in the SPI. Along with read command mode, flow control can receive multiple data bytes.

With flow control, the data transfer between the SPI master and slave is controlled based on the application requirements in terms of periodic data or demand-based data read.

The SPI master in the ADuCM4050 supports the following modes of flow control.

- Pin-based flow control, controlled by the SPI slave.
- Timer-based flow control, controlled by the SPI master.

The flow control modes are described in more detail in the Pin-Based Flow Control section and the Timer-Based Flow Control section. The mode field in the SPI flow control register (SPI_FLOW_CTL) configures the flow control mode to any one of the three modes.

Flow control mechanisms can be used only when the ADuCM4050 is configured as an SPI master.

Pin-Based Flow Control

Using a Separate RDY Pin

Some SPI slaves have a dedicated RDY pin that is connected to the RDY pin of the SPI master, which in this case is the ADuCM4050. The RDY pin is a dedicated pin (as an alternate functionality to a GPIO) for every SPI instance.

For example, the CAT64LC40 serial flash uses a dedicated RDY pin to signal the availability of data to the SPI master.

The RDY pin of the ADuCM4050 can be wired to an interrupt pin of the SPI slave in case the slave does not support a dedicated RDY pin. The slave uses the RDY pin to indicate that the acquisition and data processing is complete. The master does not provide SPI clock until it sees an active level on this pin.

The user can configure the number of bytes to be read when the RDY pin is asserted. Perform this configuration by setting the RDBURSTSZ field in the SPI flow control register (SPI_FLOW_CTL). After receiving this burst of bytes on MISO, the SPI master continues to wait for the next RDY pin assertion to receive the next set of bytes. This process is repeated until all bytes as set in the SPI count register (SPI_CNT) are received.

Using read command mode, a maximum of 16 bytes can be transmitted. This transmission is configured using the TXBYTES field of the SPI read control register (SPI_RD_CTL). The number of bytes received in one burst when using flow control is set in the RDBURSTSZ field of the SPI flow control register (SPI_FLOW_CTL). However, the total number of bytes to be received does not have an imposed maximum limit.

Using the MISO Pin

Some SPI slaves do not have a dedicated RDY pin but have a provision to reuse the MISO pin to inform the SPI master that the data is ready to be sent on MISO.

The ADuCM4050 SPI master waits for an active level transition on the MISO line and, when this is detected, reads RDBURSTSZ + 1 number of bytes and then goes back to a wait state until another active level is detected on MISO.

The polarity of the MISO/RDY pin can be configured using the POL field of the SPI flow control register (SPI_FLOW_CTL).

Timer-Based Flow Control

For slaves that do not have a dedicated pin to inform the availability of data to the master, the microcontroller uses a 16-bit timer to introduce wait states while reading data. When the timer triggers, the master reads a burst of bytes (RDBURSTSZ + 1) and then restarts the timer. The timer is clocked at the SPI clock rate (SCK), and the number of SCK cycles to wait before the timer is triggered can be set using the SPI_WAIT_TMR register. An example of this operation is shown in Figure 55.

When this scheme stalls and drives SCK for flow control, take care to ensure the last SCK edge is a sampling edge. After the stall period is over, an SCK driving edge then causes the next data transfer.

System Description

This section uses the hardware flow control mode to demonstrate how the flow control feature can contribute to power savings in a system.

The system that demonstrates this process consists of the [ADuCM4050](#) MCU and a sensor (such as an accelerometer) connected over the SPI.

To design a power efficient system, it is essential to put the core in sleep mode whenever there is no processing required. In such a system, after a sensor reading is available, the core is woken up to receive and process the data from the sensor.

The flow control and read command modes in the [ADuCM4050](#) enhances the efficiency of this process by offloading the MCU further. The system is put into Flexi mode, which keeps the core asleep and the SPI peripheral and the DMA active.

The sensor measures the data and uses the RDY pin to strobe the SPI peripheral of the data availability. There is a dedicated SPI_RDY pin (alternate functionality of a GPIO) for every SPI instance in the [ADuCM4050](#) MCU.

Without waking up the MCU, the SPI then reads the data set using the read command mode. The sensor must be capable of multibyte data transfer to use this scheme effectively. In case of an accelerometer sensor, the x-, y-, and z-axis readings are sent as six bytes over the SPI.

DMA transfers the data into an allocated memory space without CPU intervention.

The application can collect the data instantaneously after every measurement, or can collect buffered data from the slave after a configured number of bytes are collected by the sensor.

After the user defined set of bytes are collected, the SPI peripheral or the DMA can wake up the MCU to process the sensor data.

Figure 56 shows the application flow diagram for an SPI data read from an accelerometer every time an activity is detected. In sensors such as the ADXL345, the data ready interrupt reads the x-, y-, and z-axis readings in one SPI transaction. In other sensors, a FIFO configuration can be performed to store a number of samples in the sensor until the master reads the FIFO.

CONCLUSIONS

The different features of the [ADuCM4050](#) SPI, such as read command mode and flow control, make the devices ideal for use in battery-powered systems where the SPI peripheral offloads the MCU and can be independently used for data collection.

This device suitability a significant advantage in wireless sensor networks where the battery life of the sensor is critical in system design. This also serves as a building block for designing smart sensors with on-board data acquisition, as well as sensor data analytics.




Figure 55. Software Flow Control with Timer

16067-069




Figure 56. Application Flowchart

16087-070

SLEEP ON EXIT

The ARM® Cortex™-M processors are ideal for low power applications due to their balance between energy and efficiency. These processors have a feature known as sleep on exit that allows saving clock cycles and energy.

The MCU subsystem of the [ADuCM4050](#) processor is based on the ARM® Cortex™-M4F processor. The sleep on exit feature saves power when the microcontroller is sleeping and in interrupt handlers.

When sleep on exit is enabled, the processor enters directly to sleep when the ISR is finished. Interrupts are nested in case there is more than one interrupt. After the execution of these interrupts, the processor returns automatically to sleep mode.

BENEFITS

The sleep on exit feature presents some benefits in interrupt driven applications, in which the system is sleeping and it only wakes up to run interrupts.

When sleep on exit is disabled, the workflow when an interrupt arrives involves more time spent executing instructions.

The steps to perform the interrupt with sleep on exit feature disabled are as follows:

1. Wake up the processor.
2. Push all the necessary information and the current state on the stack.
3. Run the interrupt code.
4. Pop the information on the stack to restore the registers.
5. Return to sleep mode.

There are many instructions to run an interrupt. Therefore, on interrupt driven applications, the time spent on context switching is not optimal because the core pushes instructions in the stack and, subsequently, the core pops them again.

The process is simplified by enabling the sleep on exit feature. The processor immediately goes to sleep after finishing the interrupt. The device does not return to the normal thread and keeps the interrupt configuration, which avoids including push and pop tasks into the stack, saving the energy and clock cycles necessary to execute unnecessary instructions.

Figure 57 shows the flowchart when using the sleep on exit feature. The procedure for using this feature is as follows:

1. The program starts.
2. Sleep mode is invoked by a wait for interrupt (WFI) or a wait for events (WFE) instruction.
3. The system enters sleep mode.
4. The device is woken up by an interrupt or an event.
5. The system returns automatically to sleep mode when the interrupt is finished if the SLEEPONEXIT bit is set.




Figure 57. Sleep on Exit Flowchart

ENABLING THE SLEEP ON EXIT FEATURE

The ARM® Cortex™-M nested vector interrupt controller (NVIC) has a system control register with a bit field called SLEEPONEXIT. To enable the sleep on exit feature, it is only necessary to set the SLEEPONEXIT bit.

The address of the system control register is 0xE000ED10. The System Control Register in the [ADuCM4050](#) section shows this register and its bit fields in the [ADuCM4050](#) microcontroller.

SYSTEM CONTROL REGISTER IN THE ADUCM4050

Address: 0xE000ED10, Reset: 0x0000, Name: NVIC0_INTCON0




Table 26. Bit Descriptions for NVIC0_INTCON0

Bits	Bit Name	Description	Reset	Access
[15:3]	RESERVED	Reserved.	0x0	R
2	SLEEPDEEP	Deep sleep flag for hibernate mode. 0: sleep deep is not enabled. 1: sleep deep is enabled.	0x0	R/W
1	SLEEPONEXIT	Sleeps the core on exit from an ISR. 0: sleep on exit is not enabled. 1: sleep on exit is enabled.	0x0	R/W
0	RESERVED	Reserved.	0x0	R

**ESD Caution**

ESD (electrostatic discharge) sensitive device. Charged devices and circuit boards can discharge without detection. Although this product features patented or proprietary protection circuitry, damage may occur on devices subjected to high energy ESD. Therefore, proper ESD precautions should be taken to avoid performance degradation or loss of functionality.

Legal Terms and Conditions

By using the evaluation board discussed herein (together with any tools, components documentation or support materials, the "Evaluation Board"), you are agreeing to be bound by the terms and conditions set forth below ("Agreement") unless you have purchased the Evaluation Board, in which case the Analog Devices Standard Terms and Conditions of Sale shall govern. Do not use the Evaluation Board until you have read and agreed to the Agreement. Your use of the Evaluation Board shall signify your acceptance of the Agreement. This Agreement is made by and between you ("Customer") and Analog Devices, Inc. ("ADI"), with its principal place of business at One Technology Way, Norwood, MA 02062, USA. Subject to the terms and conditions of the Agreement, ADI hereby grants to Customer a free, limited, personal, temporary, non-exclusive, non-sublicensable, non-transferable license to use the Evaluation Board FOR EVALUATION PURPOSES ONLY. Customer understands and agrees that the Evaluation Board is provided for the sole and exclusive purpose referenced above, and agrees not to use the Evaluation Board for any other purpose. Furthermore, the license granted is expressly made subject to the following additional limitations: Customer shall not (i) rent, lease, display, sell, transfer, assign, sublicense, or distribute the Evaluation Board; and (ii) permit any Third Party to access the Evaluation Board. As used herein, the term "Third Party" includes any entity other than ADI, Customer, their employees, affiliates and in-house consultants. The Evaluation Board is NOT sold to Customer; all rights not expressly granted herein, including ownership of the Evaluation Board, are reserved by ADI. **CONFIDENTIALITY.** This Agreement and the Evaluation Board shall all be considered the confidential and proprietary information of ADI. Customer may not disclose or transfer any portion of the Evaluation Board to any other party for any reason. Upon discontinuation of use of the Evaluation Board or termination of this Agreement, Customer agrees to promptly return the Evaluation Board to ADI. **ADDITIONAL RESTRICTIONS.** Customer may not disassemble, decompile or reverse engineer chips on the Evaluation Board. Customer shall inform ADI of any occurred damages or any modifications or alterations it makes to the Evaluation Board, including but not limited to soldering or any other activity that affects the material content of the Evaluation Board. Modifications to the Evaluation Board must comply with applicable law, including but not limited to the RoHS Directive. **TERMINATION.** ADI may terminate this Agreement at any time upon giving written notice to Customer. Customer agrees to return to ADI the Evaluation Board at that time. **LIMITATION OF LIABILITY.** THE EVALUATION BOARD PROVIDED HEREUNDER IS PROVIDED "AS IS" AND ADI MAKES NO WARRANTIES OR REPRESENTATIONS OF ANY KIND WITH RESPECT TO IT. ADI SPECIFICALLY DISCLAIMS ANY REPRESENTATIONS, ENDORSEMENTS, GUARANTEES, OR WARRANTIES, EXPRESS OR IMPLIED, RELATED TO THE EVALUATION BOARD INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, TITLE, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS. IN NO EVENT WILL ADI AND ITS LICENSORS BE LIABLE FOR ANY INCIDENTAL, SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES RESULTING FROM CUSTOMER'S POSSESSION OR USE OF THE EVALUATION BOARD, INCLUDING BUT NOT LIMITED TO LOST PROFITS, DELAY COSTS, LABOR COSTS OR LOSS OF GOODWILL. ADI'S TOTAL LIABILITY FROM ANY AND ALL CAUSES SHALL BE LIMITED TO THE AMOUNT OF ONE HUNDRED US DOLLARS (\$100.00). **EXPORT.** Customer agrees that it will not directly or indirectly export the Evaluation Board to another country, and that it will comply with all applicable United States federal laws and regulations relating to exports. **GOVERNING LAW.** This Agreement shall be governed by and construed in accordance with the substantive laws of the Commonwealth of Massachusetts (excluding conflict of law rules). Any legal action regarding this Agreement will be heard in the state or federal courts having jurisdiction in Suffolk County, Massachusetts, and Customer hereby submits to the personal jurisdiction and venue of such courts. The United Nations Convention on Contracts for the International Sale of Goods shall not apply to this Agreement and is expressly disclaimed.