Analog and Mixed-Signal Products

Analog Applications Journal

August 2000

IMPORTANT NOTICE

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

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

Customers are responsible for their applications using TI components.

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

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

Contents

Introduction

Analog Applications Journal is a collection of analog application articles designed to give readers a basic understanding of TI products and to provide simple but practical examples for typical applications. Written not only for design engineers but also for engineering managers, technicians, system designers and marketing and sales personnel, the book emphasizes general application concepts over lengthy mathematical analyses.

These applications are not intended as “how-to” instructions for specific circuits but as examples of how devices could be used to solve specific design requirements. Readers will find tutorial information as well as practical engineering solutions on components from the following product categories:

- Data Acquisition
- Power Management
- Interface (Data Transmission)
- Amplifiers

Where applicable, readers will also find software routines and program structures. Finally, *Analog Applications Journal* includes helpful hints and rules of thumb to guide readers in preparing for their design.

Because this book is limited in size, readers should refer to more detailed technical information, which can be found on TI’s product-specific websites listed at the end of each article.

New DSP development environment includes data converter plug-ins

By Manfred Christ, *FAE Manager, Advanced Analog Products*
and Frank Walzer, *Field Application Engineer, Data Converter Software Development*

System designs that include a DSP and one or more data converters consistently increase in complexity. In addition, design engineers are asked to release new applications faster to market. To support both increasing system complexity and shorter design cycles, new development systems had to improve in code generation as well as in efficient debugging capabilities. Therefore, previous compilers were developed in close conjunction with related simulators and hardware emulators. Today all these modules have been integrated into one development system, the Code Composer Studio™ (CCS).

The CCS developments took into consideration the increasing complexity of the on-chip peripheral DSP modules such as serial and parallel ports, timer, memory, and interrupt handler. User functions, written as subroutines that configure the DSP on-chip peripherals, have been combined into libraries, thus enabling the user to execute easily the many configuration options available. In an additional step, subroutines that configure and read or write to analog-to-digital converters (ADCs) and digital-to-analog converters (DACs) have been developed for all major DSP families. The user may generate configuration data using a graphical user interface provided by a CCS plug-in.

Such a plug-in allows the designer to choose from the many options to configure a data converter via a dialog window. The dialog window provides several parameters that allow the ADC configuration to be tailored to the exact needs of the application. Once the parameters have been selected, the tool generates the necessary functions for the C-program that are required to configure the data converter. Functions to initiate the data transfer via the ADC-DSP interface and to read ADC conversion results are also generated. If a subsequent test reveals that another ADC configuration could improve the data transfer—i.e., in data throughput or in power saving—the user can return to the dialog window to select different parameters. The code generation and compilation follow with just a few mouse clicks, so that the interface then can be tested immediately for optimal performance.

Figure 1. Interfacing two THS1206 ADCs to a C6201 DSP

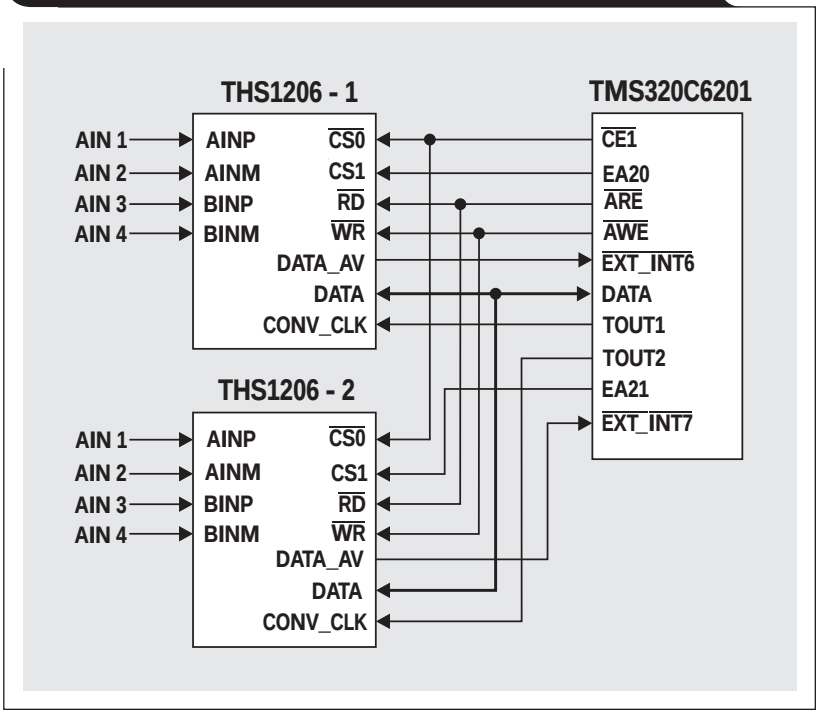

Optimized DSPs for portable and high-performance applications

Texas Instruments offers a variety of DSP architectures optimized for a wide range of applications. In addition to DSPs for motor and process controls, there are low-cost DSPs with floating-point architecture for a variety of industrial systems. To allow for high portability, the processors are often programmed in C, using the highly efficient, optimizing C-compilers.

The C5000 architecture with its many derivatives is designed for extremely low power consumption. Good examples of low-power applications are digital mobile phones such as GSM phones and PCS-1900 phones, whose standby and talk time increase with each generation. Other recent developments using the C5000 family are the Web-phones.

The C6000 DSPs offer processor performance up to several thousand MIPS (million instructions per second). Those devices are mainly used in applications requiring large processing power and high data throughput, such as compression algorithms in the audio and video segment. Typical applications are central office DSL modems as well as GSM base stations, where many data channels need to be processed simultaneously.

Continued on next page

Continued from previous page

State-of-the-art data converters offer many options

Besides the increase in DSP performance, new applications also require the increase in data converter performance. For that purpose Texas Instruments recently released the 12-bit ADC, THS1206, used for industrial control applications, and the 16-bit analog interface controller (AIC), TLV320AIC10, for telecom applications.

Modern ADCs, such as the THS1206, provide a vast amount of configuration options that enable the user to configure the ADC optimally to the system application. These options need to be initialized before the actual data transfer begins; however, during operation those options may be reprogrammed to match possible changes in system requirements.

The THS1206 samples the analog input signal at a rate of 6 MSPS (million samples per second). For applications requiring more than one analog channel, the ADC provides four single-ended, analog input channels that can be sampled simultaneously at a rate of up to 1.5 MSPS. To improve the signal-to-noise ratio (SNR), the four single-ended channels can be configured to two differential input channels. Another option is to use one differential channel and two single-ended channels.

The parallel interface of the ADC can be directly connected to the DSP without external logic. Figure 1 illustrates the interface of two THS1206 ADCs to a C6201 DSP. The DSP addresses both ADCs individually without the need for an address decoder.

The ADC contains an on-chip circular FIFO that improves the data throughput significantly. Without this FIFO, the DSP would have to read each sample separately. At a sample rate of 6 MSPS, the interrupt processing effort would increase enormously, which would not allow the DSP enough time for further signal processing. The on-chip FIFO, however, allows the DSP to read entire blocks of ADC conversion results in burst mode operation. To optimally interface the ADC to various DSPs with different processor speed, at different sample rates, the FIFO provides programmable pointers. This way the DSP is able to read the optimal number of samples after an interrupt occurs.

To optimize the ADC configuration, all available options are programmed via the two 10-bit control registers of the THS1206. The DSP needs to execute two write-cycles to define the register content with the required bit combination. Via the dialog window, Code Composer Studio generates the right bit combinations as well as the program code to initialize the ADC, based on the THS1206 plug-in.

The TLV320AIC10 offers a complete data acquisition system on one chip. The device contains ADC, DAC,

Figure 2. The "Data Converter Support" dialog window within CCS

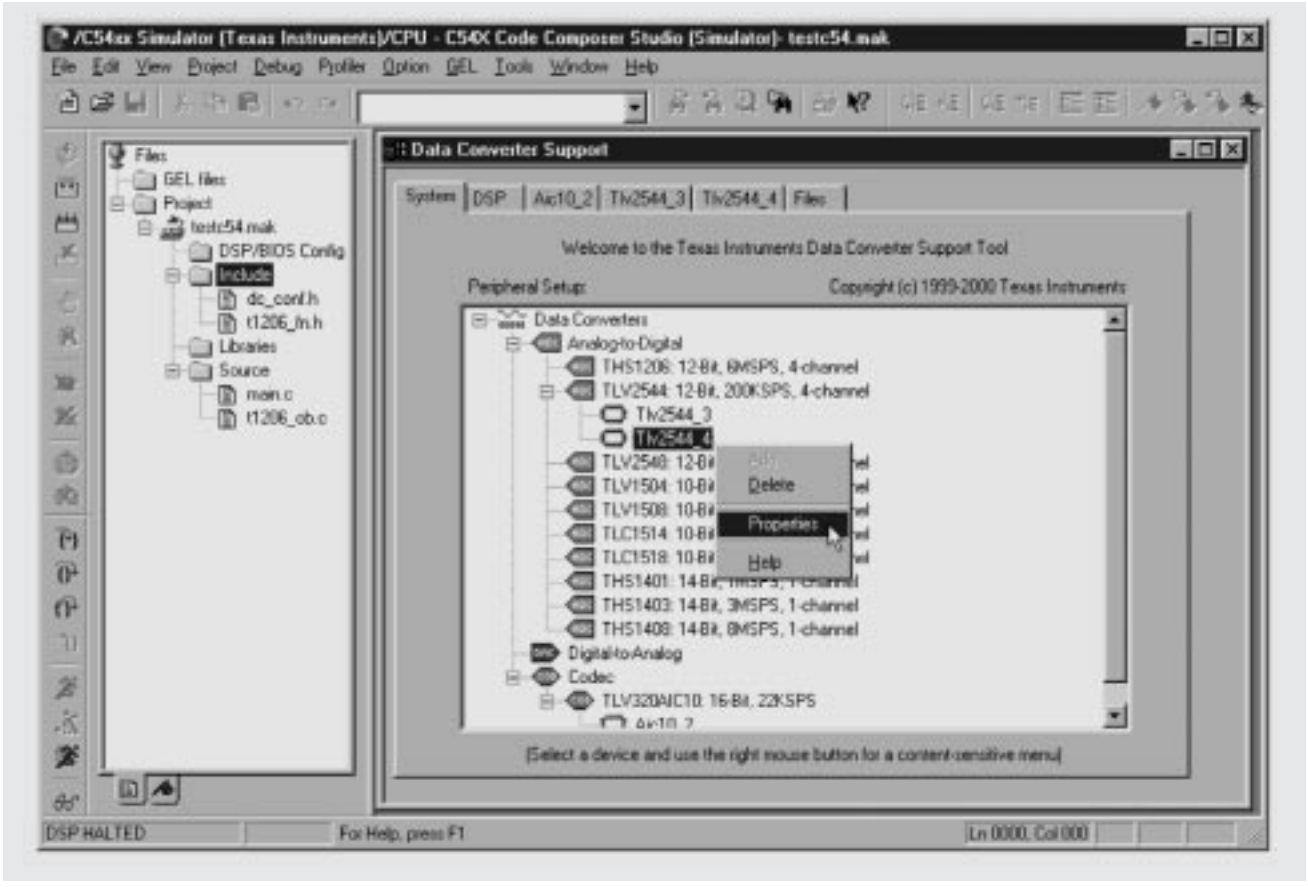

anti-aliasing filter, and many more programmable functions. The AIC10 interfaces to a DSP via a serial port. A maximum of up to eight AICs can be connected directly to one DSP serial port without external logic.

Code Composer Studio is an integrated development environment (IDE) for the latest TMS320C5000 and C6000 DSP platforms. The basic functions include C-compiler, assembler, linker, simulator, and hardware debugger. In addition, CCS provides DSP/BIOS™ technology and Real-Time Data Exchange (RTDX).

An important element is the anytime-expandable plug-in architecture. This open architecture allows third parties to integrate their specific DSP tools smoothly into CCS. System designers can choose between the plug-ins from Texas Instruments or plug-ins from third parties for a seamless implementation into their own programs. This method reduces the design cycle dramatically, since designers can focus exclusively on their essential development tasks.

Plug-ins for data converters

Texas Instruments uses the plug-in architecture to extend the development environment by gradually implementing the company's product portfolio of DSP-optimized ADCs and DACs. The goal is to simplify the designer's task dramatically, to integrate these data converters from device initialization up to the sample processing into his own algorithm.

When calling the "Data Converter Support" menu option within CCS (see Figure 2), the designer can choose from existing ADC, DAC, and AIC devices. Via the "Add" context menu in the dialog window, the desired converter is added to the system. Selecting "Properties" with the converter selected makes the configuration dialog window appear.

Figure 3 gives an example for the THS1206. This window consists of the following three segments: "Interface Settings," "ADC Settings," and "Current Control Selection." "Interface Settings," for example, includes the definition of the ADC address within the DSP memory space, the polarity of the edge of the interrupt signal, and the format of the transferred data. Within "ADC Settings," the number of analog-input channels and their operation modes is defined. In addition, the user determines the ADC conversion clock and configures the FIFO by setting the trigger level. Finally, "Current Control Selection" provides the resulting bit combination for the two control registers of the THS1206.

The following step automatically generates the source code for the ADC initialization. This code is inserted into the application project via the project manager. This way all functions are callable from the main application program.

Continued on next page

Figure 3. Various options of the THS1206 ADC for data converter setup

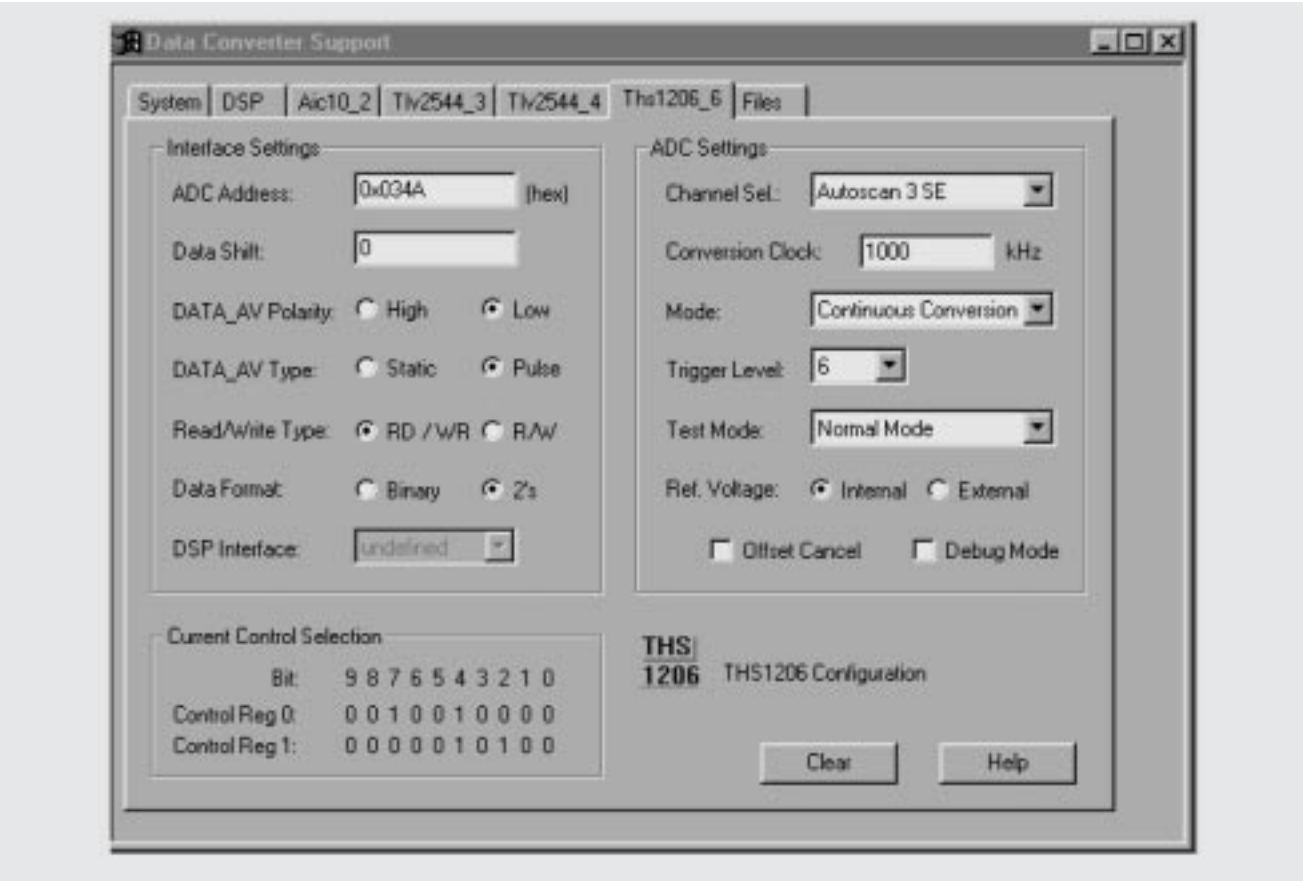

Continued from previous page

The designer can use these functions in initialization routines as well as in the actual data processing algorithms. After the source code implementation, the overall system test can begin. If the user discovers that different settings are required, he returns to the "Data Converter Support" dialog window and selects new parameters.

Future outlook

The current version of CCS already includes support for 11 data converters. Further devices currently under development will be integrated gradually into the plug-in. An installation program, updated with the latest plug-in, is available for downloads from TI's Web site at www.ti.com/sc/dcplug-in. This program is downward-compatible with previous CCS versions. Every update of the CCS CD-ROM includes the latest plug-in version. Program examples for some of the DSP DSKs and data converter EVMs are also available. Future standard software interfaces for ADCs and DACs are currently under development.

References

For more information related to this article, you can download an Acrobat Reader file at www-s.ti.com/sc/techlit/*litnumber* and replace "*litnumber*" with the **TI Lit. #** for the materials listed below.

Document Title	TI Lit. #
1. Code Composer Studio User's Guide	spru328
2. Data Acquisition Data Book	—

Related Web sites

www.dataconverter.com

http://dspvillage.ti.com.docs/ccstudio/ccstudiohome.jhtml

www.ti.com/sc/select

www.ti.com/sc/docs/products/analog/ths1206.html

www.ti.com/sc/docs/products/analog/tlv320aic10.html

www.ti.com/sc/docs/products/dsp/tms320c6201.html

www.ti.com/sc/docs/products/dsp/c5000/bench.htm

www.ti.com/sc/docs/products/dsp/c6000/apps.htm

Higher data throughput for DSP analog-to-digital converters

By Heinz-Peter Beckemeyer

Field Application & System Engineer, Data Converters

The processing power of today's digital signal processors (DSPs) goes into the range of several MIPS (million instructions per second). However, these high data rates can be achieved only if algorithms are executed from the internal memory. The maximum communication speed with external peripherals, such as analog-to-digital or digital-to-analog converters, is highly limited due to external bandwidth.

At conversion rates above 1 MSPS (million samples per second), it is frequently difficult to collect the data from the analog-to-digital converter (ADC) in time to process it further.

Data converters with such a high conversion rate most likely have a parallel interface to the parallel data bus of a processor, since the speed on serial interfaces is further limited. The reading of data from such an ADC by a processor normally takes place by means of an interrupt service routine (ISR). The maximum speed, however, is limited, as every processor has a latency time (the time required for first possible access within the ISR). The latency time can add up to some hundred nanoseconds depending on the processor type and its frequency. The resulting speed is thereby limited to 1 to 2 MHz. From this arises the problem of capturing and processing data that is controlled via an interrupt from a fast ADC.

This article suggests a solution that increases the data throughput of an ADC into a DSP by a factor of 5 to 10. This solution is already integrated in one of the latest ADCs, THS1206, developed by Texas Instruments.

Conversion and transfer of data

Fast analog-to-digital (flash/pipeline) converters typically provide a digital conversion value with every clock cycle of the sampling clock.

It frequently becomes a challenge for the DSP to collect the data of these converters and process it at this higher speed. For example, a data converter with a sample rate of 6 MSPS brings a new conversion value to the digital output every 167 ns.

The maximum data throughput is a limitation because of the latency time already mentioned. The latency time of the TMS320C542 DSP with a frequency of 40 MHz is around 400 ns. This latency time is lost every time the DSP jumps into the ISR to read a conversion value.

Higher data throughput by using a FIFO

A higher data throughput can be achieved when the DSP can read a block of conversion values instead of only one value with every jump into the ISR. In this case the DSP loses the latency time only once per block to be read.

Figure 1 shows an example where the DSP always reads 8 conversion values in a block within an ISR, therefore losing the latency time only once per ISR.

Figure 1. Reading a block of data

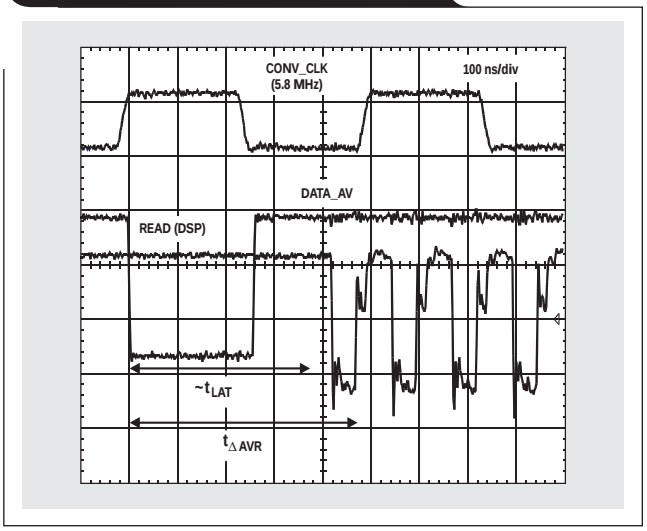

The latency time can be seen in Figure 1. It is the time from the falling edge of the signal DATA_AV (data available) to the first reading of the DSP.

The reading of the data in a block can be achieved by using a FIFO architecture. The THS1206, which is a 12-bit and 6-MSPS ADC, consists of a FIFO organized as a circular buffer to optimize the efficiency between ADC and DSP.

Figure 2 shows the block diagram of the THS1206. The main feature on the analog side is the provision of four analog input channels, which are switched simultaneously from the sample mode to the hold mode (simultaneous sample and hold). This feature is very important in modern control applications, radar and communications systems. The FIFO is capable of storing up to 16 conversion values and is located between the ADC core and the digital output of the THS1206.

Digital interface of the THS1206

When a digital interface for a DSP is evaluated, care should be taken to ensure that the THS1206 can be connected to different kinds of processors without external logic. To this end, one important criterion is the programmable reading and writing input. Processors today have either a combined reading/writing (R/$\overline{W}$) output or separate outputs for reading ($\overline{RD}$) and writing ($\overline{WR}$). The THS1206 can be connected to both types of processor without external logic. As can be seen in Figure 2, the THS1206 possesses a reading ($\overline{RD}$) input and a writing ($\overline{WR}$) input. The $\overline{WR}$ input can be reprogrammed on startup to a combined

Continued on next page

Figure 2. Block diagram of the THS1206

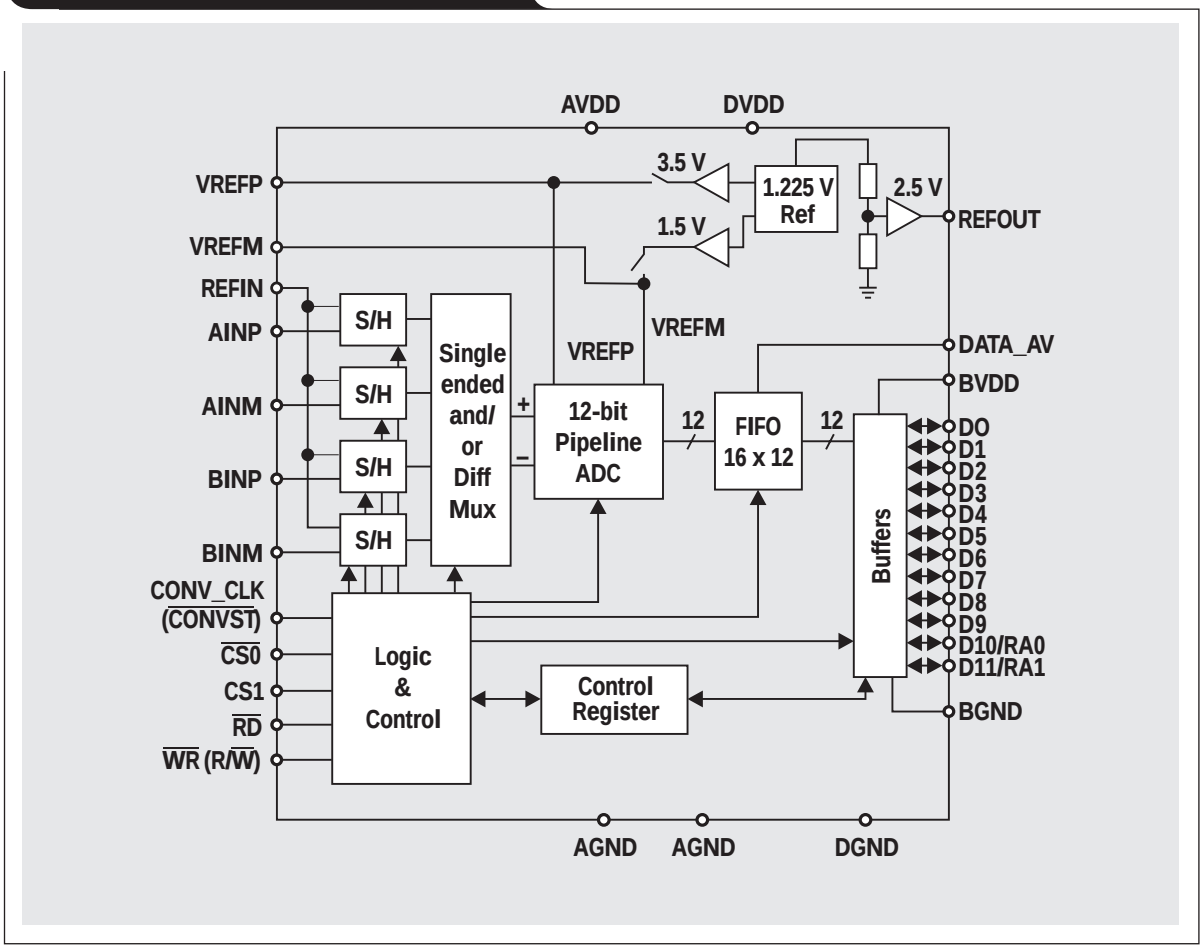

Continued from previous page

reading/writing input by a corresponding programming of the internal control register. With the examples of the two DSPs, TMS320C54x and TMS320C6201, the different possibilities are shown.

Interface of the THS1206 to the TMS320C54x

Figure 3 shows an example in which a THS1206 is connected to the DSP TMS320C54x. The TMS320C54x controls the reading and writing via the combined output R/$\overline{W}$. The R/$\overline{W}$ input of the THS1206 is programmed to this end as a combined reading and writing input. The $\overline{RD}$ input is switched to inactive in this operating mode. Addressing the THS1206 occurs via the two inputs $\overline{CS0}$ and CS1. The selection of the I/O peripherals is controlled via the output $\overline{IOSTRB}$ of the TMS320C54x, and addressing takes place via the address line A14. The converted data is read by TMS320C54x, controlled via an interrupt ($\overline{INTX}$).

Figure 3. Interface of the THS1206 to the TMS320C54x

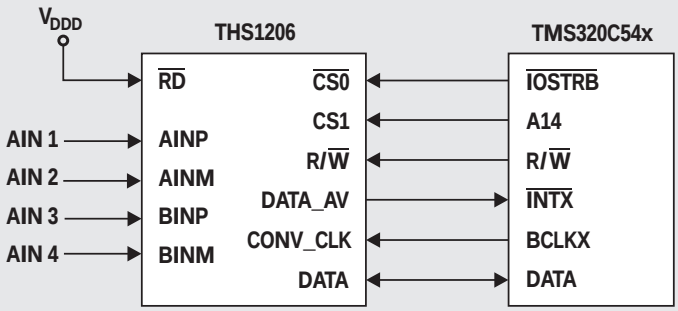

Interface of the THS1206 to the TMS320C6201

Figure 4 shows an example in which two THS1206 ADCs are connected to the TMS320C6201 DSP. The TMS320C6201 controls the reading via the output $\overline{ARE}$ and the writing via the output $\overline{AWE}$. The selection of the THS1206 ADC is undertaken via the control output $\overline{CE1}$. The addressing of both ADCs occurs via the address connectors EA20 and EA21. An external address decoder therefore is not required. The converted data is at any time controlled via an interrupt ($\overline{EXT\_INT6}$, $\overline{EXT\_INT7}$) taken over by TMS320C6201.

Figure 4. Interface of the THS1206 to the TMS320C6201

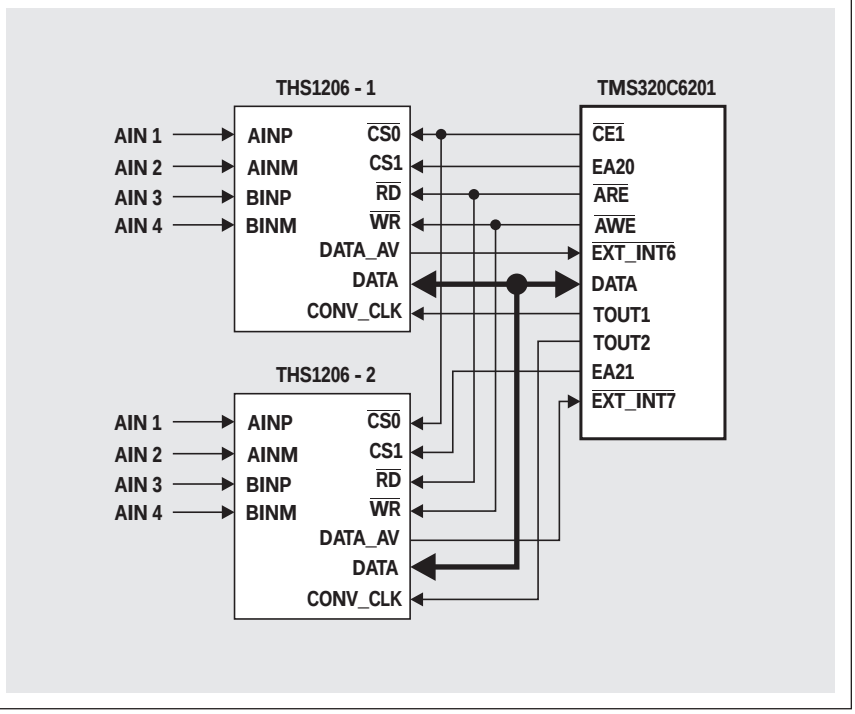

Integrated FIFO—circular buffer

As previously mentioned, the new THS1206 ADC has an on-chip FIFO, implemented as a flexible circular buffer to increase the maximum data throughput between ADC and DSP. With every falling edge of the sampling clock, a new conversion is initiated and a converted value is written into the FIFO. The circular buffer integrated in the THS1206 can store up to 16 conversion values. With a conversion rate of 6 MHz and a programmed FIFO depth of 8, the THS1206 stores 8 values in the FIFO within 8 x 167 ns = 1.34 µs. Only then does the ADC indicate to the DSP, with the digital control signal DATA_AV, that a block of data is ready to be read. This gives more flexibility to the processor and frees up resources for other activities.

The FIFO is configured as a flexible circular buffer. In principle it allows an overwriting of data that hasn't been read. This can be of interest for applications where data has to be sampled only in specific time frames. An example of this is the fast Fourier transformation (FFT), where typically a block of data has to be read (1024, 2048, 4096, ...).

However, the data that is read always should be the most recent, which is ensured by this specific circular buffer architecture. Overwriting is prevented in the relevant time frames by reading fast enough from the processor to prevent overlap.

Figure 5 shows the architecture of the integrated FIFO, organized as a circular buffer. The reading from and writing to the FIFO is controlled entirely internally and can occur asynchronously. The conversion values are written automatically to the FIFO. In order to control the reading and writing, a read pointer and write pointer are used. These pointers always point to the cell in the FIFO where data has been written and to the cell that should be read next time.

This FIFO is designed specifically to allow multiple analog input channels for the ADC. In this case, the converted values are written into the FIFO in a predefined sequence (Autoscan Mode). Therefore, the channel information is always maintained for the processor and the higher data throughput can also be achieved for a data converter with multiple inputs.

Figure 5. FIFO organized as a circular buffer

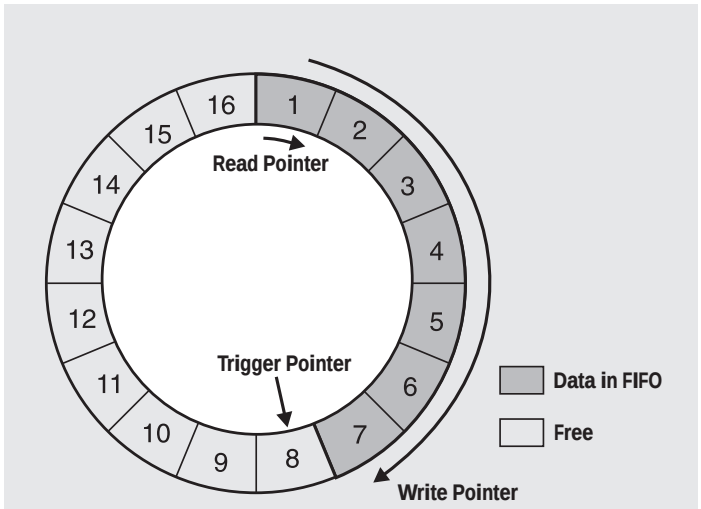

Continued on next page

Continued from previous page

It is possible to select a specific storage depth (trigger level). When this level is reached, the converter (THS1206) informs the connected processor via the digital output DATA_AV (data available) that a block of converted values is ready to be transferred. The block size to be read is always equal to the setting of the storage depth (trigger level). A further pointer (trigger pointer) controls the signal DATA_AV dependent on the storage depth.

Figure 5 gives an example with a selected storage depth of 8. Initially, the trigger pointer points to position 8. The read pointer points to position 1, and the write pointer initially begins to write into location 1 and is incremented with every conversion value. In the example, the write pointer already points to position 7. The next increment of the write pointer reaches the trigger pointer, which means that the storage depth is reached (position 8). Therefore, the control signal DATA_AV becomes active and tells the processor that a block of data (8 values) waits to be read. The selectable storage depth allows variable settings for different processors and specific applications.

Figure 6. Increased data throughput for the TMS320C542

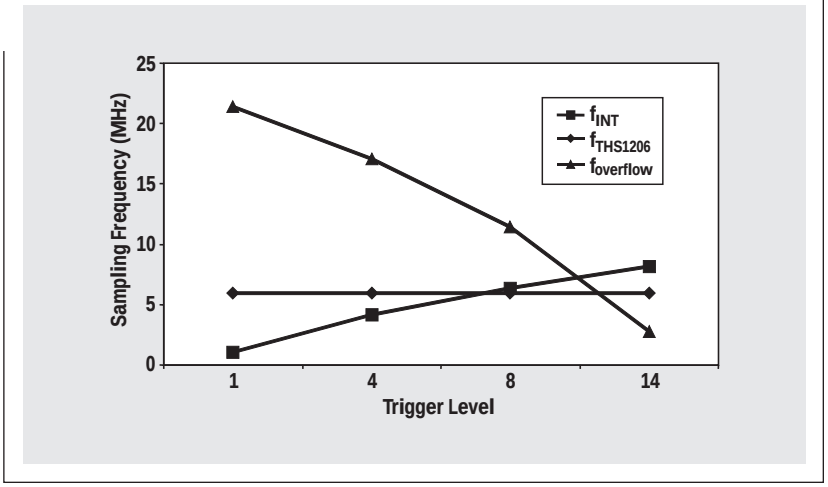

Efficiency of the FIFO

As previously explained, the conversion values of the ADC can be read in a block from the processor. This optimizes the interface between the ADC and the DSP, which frees up the processor and speeds up the data transfer.

To analyze the maximum data throughput between the processor and ADC, the latency time of the processor, the length of the ISR, and the storage depth of the FIFO are essential. The required parameters for an analysis of the maximum data throughput are:

T_F Depth of the FIFOs
T_L Trigger level
$t_{\Delta AVR}$ DATA_AV active to first read
t_{LAT} DATA_AV active to first instruction within the ISR
t_{ISR} Length of the ISR

The latency time is defined by the hardware and software of the DSP. The length of the interrupt service routine results from the latency time and the time required to read a block of data from the FIFO. The data throughput goes up with an increasing block size that can be read within one ISR. This is valid up to a specific limit of the FIFO, determined by the depth of the FIFO. As soon as the storage depth (trigger level) reaches the maximum depth of the FIFO, old data might be overwritten. Therefore, for every processor and specific application, the ideal choice of all parameters exists.

The maximum data throughput can be derived as follows. Equation 1 determines the maximum data throughput at a given trigger level (T_L) and depth of the FIFO (T_F) without overwriting of old data. Equation 2 determines the ability of the processor to capture the data fast enough at a given trigger level (T_L) and latency time (t_{LAT}).

$$f_{overflow} = \frac{T_F - T_L}{t_{\Delta AVR}} \quad (1)$$

$$f_{INT} = \frac{T_L}{t_{LAT} + t_{ISR}} \quad (2)$$

Figure 6 proves the efficiency of the FIFO. In this example, the THS1206 is connected to the TMS320C542 with a 40-MHz crystal. The figure shows the maximum sampling rate versus the selected storage depth (trigger level). The speed limitation caused by the length of the ISR (f_{INT}), the risk of overwriting unread data ($f_{overflow}$), and the maximum sampling speed of the ADC ($f_{THS1206}$) are shown. The highest data throughput at a specific trigger level is determined by the slowest of these three parameters. In this specific example, the maximum speed of the THS1206 can be achieved only if a trigger level of 8 is chosen for a 16-word-deep FIFO. At a trigger level below 8, the processor is unable to achieve the high data throughput because of the latency time. Overwriting of old data becomes critical with a trigger level higher than 8.

The maximum data throughput of the THS1206 into the TMS320C54x also is shown in Table 1.

Table 1. Data throughput of the TMS320C542-40

T_L	1	4	8	14
t_{ISR}	475 ns	550 ns	850 ns	1300 ns
f_{device}	6 MHz	6 MHz	6 MHz*	6 MHz
$f_{overflow}$	21.4 MHz	17.1 MHz	11.4 MHz	2.8 MHz*
f_{ISR}	1.1 MHz*	4.2 MHz*	6.4 MHz	8.2 MHz

*Maximum data throughput dependent on the selected trigger level

Figure 7 shows a second example of the FIFO's efficiency. In this case the THS1206 with integrated FIFO is connected to the TMS320C6701 with a 100-MHz crystal. The figure shows the maximum sampling rate versus the selected storage depth (trigger level). The speed limitation caused by the length of the ISR (f_{INT}), the risk of overwriting unread data ($f_{overflow}$), and the maximum sampling speed of the ADC ($f_{THS1206}$) are shown. The highest data throughput at a specific trigger level is determined by the slowest of these three parameters. In this specific example, the maximum speed of the THS1206 can be achieved if a trigger level higher than 1 is chosen for a 16-word-deep FIFO. Even a trigger level of 14 is possible in this example. The processor is fast enough to react and to prevent an overwriting.

The maximum data throughput of the THS1206 into the TMS320C6701 also is shown in Table 2.

Figure 7. Increased data throughput for the TMS320C6701

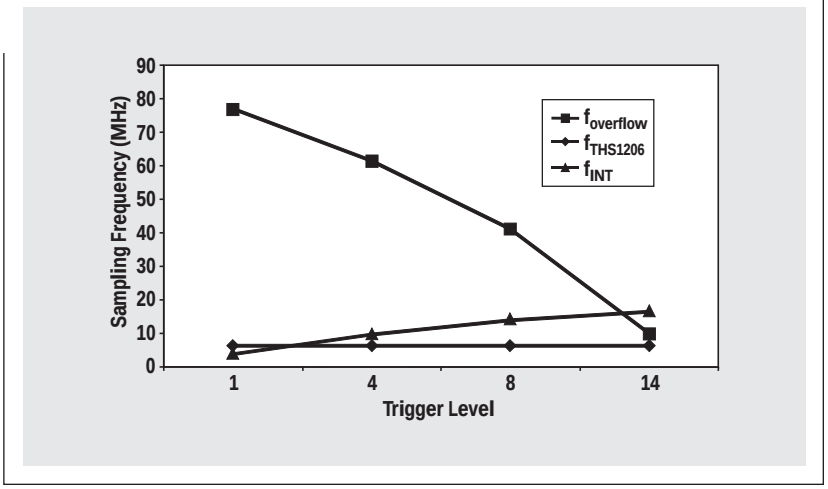

The explanations and examples used in this article show that the way data converters and DSPs transfer data will have a significant effect on the overall performance of a system. System designers will find that the speed of these data transfers can be optimized when the data converters and DSPs have features and capabilities that complement each other. This requires an intimate knowledge of both DSP and data converter technology.

Texas Instruments provides a family of ADCs with FIFO. These are the THS10064 (10 bits, 6 MSPS, 4 analog inputs), THS12082 (12 bits, 8 MSPS, 2 analog inputs), and the THS10082 (10 bits, 8 MSPS, 2 analog inputs).

Table 2. Data throughput of the TMS320C6701-100

T_L	1	4	8	14
t_{ISR}	130 ns	280 ns	480 ns	780 ns
f_{device}	6 MHz	6 MHz*	6 MHz*	6 MHz*
$f_{overflow}$	76.9 MHz	61.5 MHz	41.0 MHz	10.2 MHz
f_{ISR}	3.9 MHz*	9.6 MHz	13.1 MHz	15.4 MHz

*Maximum data throughput dependent on the selected trigger level

References

For more information related to this article, you can download an Acrobat Reader file at www-s.ti.com/sc/techlit/*litnumber* and replace *"litnumber"* with the **TI Lit. #** for the materials listed below.

Document Title	TI Lit. #
1. Data Book Data Acquisition Circuits (1998)	—
2. Data Converter Selection Guide (1998)	—
3. Mixed-Signal & Analog Products CD-ROM (1999)	—
4. Designing with the THS1206 High-Speed Data Converter	slaa094

Related Web sites

www.dataconverter.com
www.ti.com/sc/select
www.ti.com/sc/1206
www.ti.com/sc/docs/products/dsp/tms320c542.html
www.ti.com/sc/docs/products/dsp/tms320c54x.html
www.ti.com/sc/docs/products/dsp/tms320c6201.html
www.ti.com/sc/docs/products/analog/ths10064.html
www.ti.com/sc/docs/products/analog/ths12082.html
www.ti.com/sc/docs/products/analog/ths10082.html

Efficiently interfacing serial data converters to high-speed DSPs

By Lijoy Philipose
Application Specialist, Data Converters

Introduction

Efficiently interfacing serial analog-to-digital data converters to TMS320C54xx digital signal processors (DSPs) can seem like a daunting task. Typically, these data converters are ten times slower than the high-speed DSPs interfaced to them. In high-performance, time-critical applications, asking the DSP to service these slower off-chip devices is a waste of processing power. The desirable solution is to transfer the task to an off-chip peripheral device. The direct memory access (DMA) controller used in conjunction with the multi-channel buffered serial port (McBSP) peripheral provides such a solution. This system will interface to the slower data converter by reading and storing samples into memory for DSP processing. The goal here is to introduce the McBSP and the DMA-enhanced peripherals and to provide an outline for interfacing serial data converters to the McBSP of the C5402 DSK. This article highlights how to use the McBSP and DMA to read samples from a serial analog-to-digital converter (ADC), specifically the TLV2548 ADC. The advantage of this method is that it frees the DSP for work more deserving of its power and speed.

Figure 1. McBSP sub-address scheme

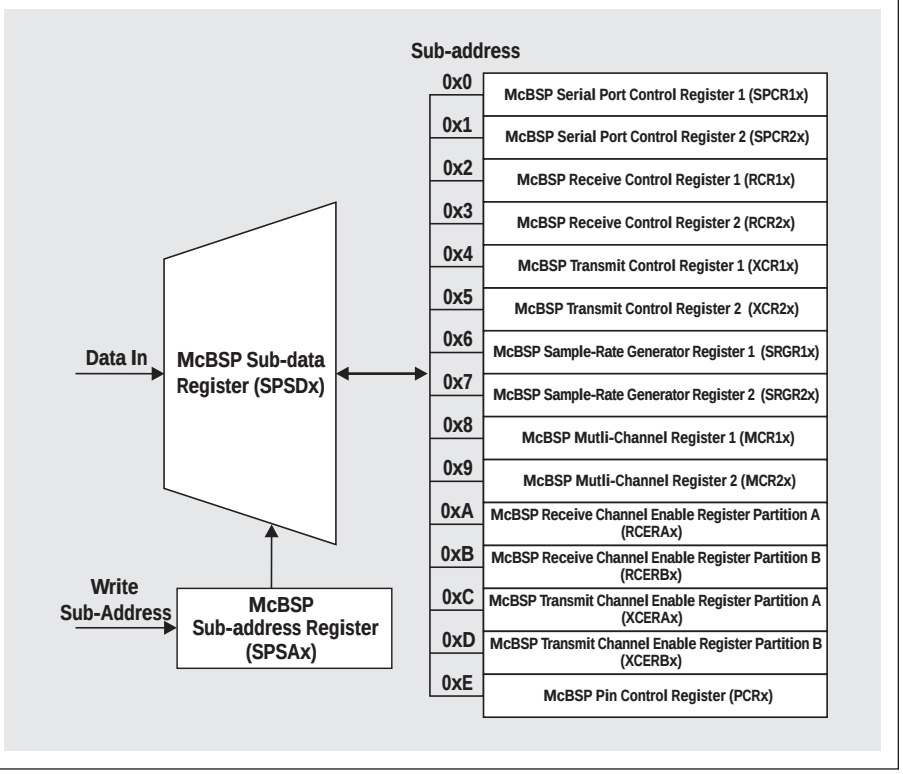

McBSP

The multi-channel buffered serial port is a superset of the standard serial ports found on TI's DSPs. The McBSP not only has features found on previous serial port interfaces but also can direct interfacing to TI/E1 framers, IOM-2 compliant devices, MVIP-switching compatible and ST-BUS compliant devices, AC97 compliant devices, IIS compliant devices, and SPI devices. It provides a wide selection of transmit/receive data sizes, µ-law and A-law companding, programmable polarity for both frame synchronization and data clocks, and a highly programmable internal clock and frame generation. These features and programming requirements are described in Reference 6.

Let's look at how the McBSP registers are accessed. The McBSP registers are memory-mapped using a register sub-addressing scheme. Figure 1 shows a visual representation of this scheme. Register sub-addressing involves multiplexing a set of registers to a single location in the memory map. A sub-bank address register is used to control the multiplexer. A sub-data register (SPSDx) is used to read or write data to the desired sub-addressed register. To access a specific sub-addressed register, the register's sub-address location is written into the sub-address register (SPSAx). This directs the multiplexer to connect to the desired physical location in memory. When a write access occurs, data written to the sub-data register is moved to the embedded data register specified in the sub-address register. Similarly, for a read access, the contents of the register specified in the sub-address register are moved to the sub-data register.

Take, for example, McBSP0. The sub-data register (SPSD) is at location 0x039, and the sub-address register (SPSAx) is at location 0x038 in physical memory. The following sample code writes 0x000 to Serial Port Control Register 1 of McBSP0.

```
SPSA0        .set   038h   ;McBSP0 Sub-Address Register
SPSD0        .set   039h   ;McBSP0 Sub-Data Register
SPCR10_SUB   .set   000h   ;McBSP0 Serial Port Control
                             Register 1 sub-address

mmr(#SPSA0) = #SPCR10_SUB
mmr(#SPSD0) = #0000h
```

There are 16 registers associated with the McBSP. Interfacing a single TLV2548 ADC to the McBSP requires the proper configuration of only the following 9 registers. The serial port is configured by properly initializing the following registers:

1. *Serial Port Control Register 1 (SPCR1)* contains the McBSP receiver status bits and the main switch to enable or disable the receiver. This register includes the clock stop mode bit, which sets the serial port for various clocking modes for SPI and non-SPI schemes. Also included in SPCR1 is the ABIS mode bit and receiver interrupt mode bit.
2. *Serial Port Control Register 2 (SPCR2)* contains the McBSP transmitter status bits and the main switch to enable or disable the transmitter. This register also contains the bits to reset the frame-sync generator and the sample-rate generator.
3. The *Pin Control Register (PCR)* contains the bits to configure the McBSP pin as inputs or outputs during normal serial port operation. This is used to reconfigure the serial port pins as general-purpose inputs or outputs when the receiver or transmitter is disabled. PCR configures the transmitter and receiver clock and frame sync modes. For example, these bits determine whether CLKX/R and FSX/R are input/output pins and what their polarity is.
4. *Receive Control Register 1 (RCR1)* contains the bits to configure various options of the receiver. The value of this register determines the receiver word size (between 8 and 32 bits) and the number of words per frame (1 to 128) expected per receiver event.
5. *Receive Control Register 2 (RCR2)* determines the size of the word received and the bit delay after the frame-sync pulse. This register configuration plays an essential role if transfers greater than 16 bits and multiple phases are necessary. In this application the register bit of interest is the data-bit delay. The RCR2 register bits also select between µ-law, A-law companding, or whether the MSB or the LSB is transferred first for non-companding 8-bit transfers.
6. *Transmit Control Register 1 (XCR1)* contains the bits that determine the transmit word length and frame size. The transfer can be 8 to 32 bits wide and anywhere from 1 to 128 words long.
7. *Transmit Control Register 2 (XCR2)* contains the bits that determine the transmit data delay and select between µ-law, A-law companding, or whether the MSB or the LSB is transferred first for non-companding 8-bit transfers.

8. & 9. *Sample-Rate Generator Register 1 (SRGR1)* and *Sample-Rate Generator Register 2 (SRGR2)* control the sample-rate generator. The sample-rate generator is composed of a three-stage clock divider that allows programmable data clocks (CLKG) and framing signals (FSG). These two McBSP internal signals can be programmed to drive receive/transmit clock (CLKR/X) and receive/transmit data framing (FSR/X). Sample-rate generator registers (SRGR[1,2]) control the operation of the various features of the sample-rate generator. These registers are used to control the width of the frame-sync pulse and to determine whether frame-sync is an external input driven by a sample-rate generator or a signal indicating that a data copy from DXR(1,2) to XSR(1,2) has been made. These registers control whether the sample-rate generator clock is derived from the CPU clock or CLKS pin, and by what value to divide the CPU clock to produce the desired serial clock (CLKX/R).

Describing all 9 registers in detail is not appropriate in this forum. Detailed descriptions of all the register bits are provided in Reference 6.

The McBSP registers are initialized to the same value regardless of whether the CPU or the DMA acts as interface controller. Table 1 is a summary of the McBSP register values typically used in most serial interfaces. These same settings can be used for most serial data converters that provide less than 16 bits of resolution.

Continued on next page

Table 1. McBSP register settings

McBSP1 ADDRESS	McBSP1 SUB-ADDRESS	ACRONYM	REGISTER INITIALIZED	COMMENT
0041		DRR11		Receiver data register.
0043		DXR11		Transmit data register.
0048		SPSA1		McBSP1 sub-addressing register.
0049	0x0000	SPCR11	0x0001	While configuring McBSP transmitter-related registers, the LSB bit must be 0. This disables the transmitter.
	0x0001	SPCR21	0x02C1	While configuring McBSP receiver-related registers, the LSB bit must be 0. This disables the receiver.
	0x0002	RCR11	0x0040	Selects one 16-bit word transfer per frame.
	0x0003	RCR21	0x0001	Set 1-bit delay on receiver. Receiver assumes first MSB bit to arrive during clock cycle following FSR pulse.
	0x0004	XCR11	0x0040	Selects one 16-bit word transfer per frame.
	0x0005	XCR21	0x0001	Transmitter shifts out data immediately following falling edge of FSX.
	0x0006	SRGR11	0x0009	Assuming 100-MHz CPU clock, the CLKX has a frequency of 10 MHz. CLKX = CPU clock/(CLKGDV + 1), where CLKGDV = SRGR1(7,0).
	0x0007	SRGR21	0x2000	The sample-rate generator clock is derived from the CPU clock.
	0x000E	PCR1	0x0A00	FSX is determined from sample-rate generator frame-synchronization mode bit SRGR2.FSGM. CLKX output is driven by the sample-rate generator.

Continued from previous page

Direct memory access (DMA)

The DMA controller uses the same sub-addressing scheme as the McBSP, so the same procedure used in writing and reading from McBSP registers must be used here. The DMA provides one additional register, which makes programming the registers easier. Figure 2 shows two registers at the input to the multiplexer. DMSDI is a sub-bank access register (DMSA) that increments the sub-bank address register after each read/write, whereas the DMSDN register does not increment the sub-bank address register after read/write operations. The obvious advantage to using this DMSDI is that the user no longer has to change the contents of DMSA to point to the next register. For example, programming DMA channel 0 registers involves writing 0x00 to DMSA. The first value written to DMSDI is moved to DMSRC0, the second to DMDST0, the third to DMCTR0, etc. The DMSDN register can be used to write to the DMA sub-register if that particular register is the only one to be modified. Five channel-specific registers and the Channel Priority and Enable Control Register (DMPREC) need to be configured for each DMA channel used. They are outlined as follows:

Figure 2. DMA sub-address scheme

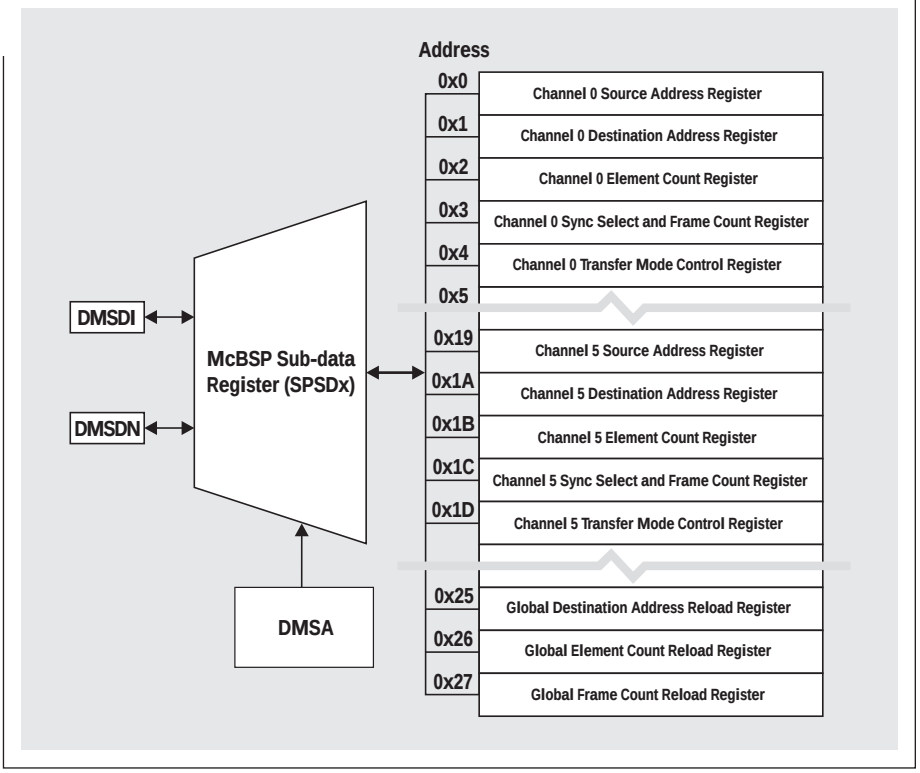

1. *Source (DMSRCn) and Destination (DMDSTn) Registers* (where n is the DMA channel). The DMSRC register stores the address of the data to be read out or moved from. Likewise, the DMDST register contains the address to where the data will be written in memory.
2. The *Element Count Register (DMCTRn)* is a 16-bit counter that keeps track of the number of DMA transfers to be completed. This register is always initialized to one less than the number of elements to be stored.
3. The *DMA Sync Event and Frame Count (DMSFCn) Register* controls three services: 1) the synchronization event used to trigger a DMA transfer; 2) the word size for each transfer, specified as either 16-bit or 32-bit words; and 3) the number of frames to be transferred. A frame can be specified from as low as 1 up to 256. For example, if only one frame is desired, this register field should be written as zero (the desired value minus one).
4. The *Transfer Mode Control Register (DMMCRn)* controls the channel transfer mode. This register determines whether the source/destination address will be post-increment or post-decrement after each transfer, whether the channel is operating in auto-buffering mode (ABU) or multi-frame mode, when the DMA will interrupt the DSP, and what address space the source/destination addresses.
5. The *Channel Priority and Enable Control Register (DMPREC)* controls the high-level function of the DMA channels. Due to the limited number of interrupts available in the C54xx family, some DMA interrupts are multiplexed with other peripheral interrupts; the register bits in this register determine which interrupts are assigned to the Interrupt Flag Register. The DMPREC also sets the priority given to each channel, either low or high. Channels with high priority are serviced before those with low priority.

The DMA register values used in this article are presented in Table 2.

Now that we have a basic understanding of the McBSP and DMA peripherals, let's look at the TLV2548 analog-to-digital converter.

TLV2548

This ADC features 8 analog input channels, an 8-level FIFO, and various conversion modes. The TLV2548 12-bit ADC is interfaced easily to the C5402 DSK. The hardware connections to the McBSP1 are shown in Figure 3. The DSP/McBSP drives serial lines CLKX and FSX, and accepts as input CLKR and FSR. The serial lines CLKX and CLKR, as well as FSX and FSR, are tied together on the EVM provided by Texas Instruments. The TLV2548 interrupt line is tied to INT3# of the DSP. Note that, since the FSX and FSR signals are identical, whenever a transmit operation occurs, a receiver transfer operation also occurs.

The TLV2548 offers four conversion modes, but only one is of interest—the single-shot conversion mode. In

Table 2. DMA register settings

REGISTER NAME	REGISTER VALUE	COMMENT
DMSRC0	&DRR11	*DMA Channel 0: Source Address Register* Source memory-mapped address to read from
DMDST0	&DataTable	*DMA Channel 0: Destination Address Rregister* Destination address to store sample to
DMCTR0	NSAMPLES-1	*DMA Channel 0: Element Count Register* Number of samples to store minus one
DMSFC0	0x5000	*DMA Channel 0: Sync Select and Frame Count Register* Synchronization transfers with McBSP1 receive event. When McBSP REVT occurs, CH0 reads data out of DRR1; then CH1 moves channel command word to DXR1.
DMMCR0	0xC004	*DMA Channel 0: Transfer Mode Control Register* Channel-context registers are re-initialized upon completion of block transfer. DMAC0 interrupt is generated after block transfer. No modification is made to Source Address Register after each transfer. Destination Address Register is post-incremented after transfer.
DMSRC1	&ADC_Cmd	*DMA Channel 1: Source Address Register* Address in memory where conversion command word is located.
DMDST1	&DXR11	*DMA Channel 1: Destination Address Register* Destination address to store conversion command—i.e., memory-mapped address of McBSP1 transmit data register.
DMCTR1	NSAMPLES -1	*DMA Channel 1: Element Count Register* Number of times to transfer command word to transmitter minus one.
DMSFC1	0xE000	*DMA Channel 1: Sync Select and Frame Count Register* Synchronization transfers with INT3# receive event. When INT3# occurs, CH0 reads data out of DRR1; then CH1 moves channel command word to DXR1.
DMMCR1	0x8000	*DMA Channel 1: Transfer Mode Control Register* The contents of the registers are re-initialized upon completion of block transfer. No modification is made to Source Address Register after each transfer. No modification is made to Destination Address Register after each transfer.
DMAPREC	0x0103	*DMA Priority and Enable Control Register* Enable DMA CH0 and CH1. CH0 has high priority. CH0 reads data out of the DRR1 and stores it before CH1 triggers another conversion cycle.

this mode, the digital code from the previous conversion must be read out before triggering another conversion cycle. This mode requires the CPU to service it often, preventing the CPU from performing other tasks. For most applications, it is undesirable to tie up an ultra-fast DSP by having it baby-sit a slow off-chip device. The solution is to use an off-chip peripheral, the DMA controller. This controller transfers data between regions in memory without intervention of the CPU. One DMA channel can be programmed to take a command word (for ADC) from memory and place it in the McBSP transmit data register, while another independent DMA channel reads the data out of the McBSP receive data register and stores it at a specified memory location (data table). Figure 3 shows the block diagram of this method. The source register of DMA Channel 1 points to the address of the next conversion command word. This command word tells an ADC which mode to enter, or what channel to convert from next. The destination register contains the address of the McBSP transmit data register. Likewise, the DMA Channel 0 source register contains the address of the McBSP receive data register, and its destination register contains the memory address to store the ADC digital code. Each DMA channel's transfer mode control register contains the bits that determine whether the source/destination addresses decrement, increment, or remain unchanged. If a number of samples is desired from only one channel, then the DMA channel transfer mode register is set so that no modification is made to the source and destination registers. The receiver DMA Channel 0 Transfer Mode Control Register is configured so the destination register is incremented after every transfer. This ensures that every read sample from the ADC is stored in a separate location in memory.

Figure 3. C5402 DMA interfacing to TLV2548

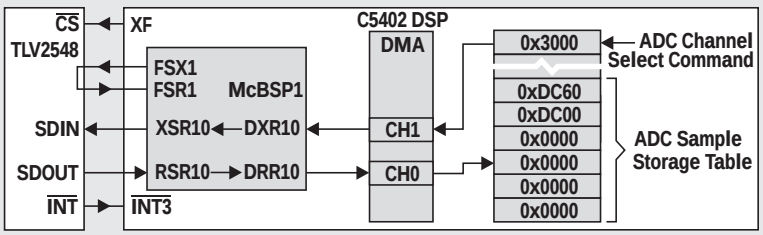

Continued on next page

Continued from previous page

The software flowchart for this interface method is presented in Figure 4. The CPU begins by initializing the respective DSP, McBSP, and DMA registers. Afterward, the only two remaining tasks for the CPU to perform are to send the configuration word and conversion channel command word to the ADC. Once the command word is placed in the data transmit register, it is automatically transferred through the transmit shift transfer onto the DX serial line. The McBSP sends out an FSX pulse immediately before the MSB bit is shifted out on the DX line. During this time, the ADC begins to shift out data on the SDO line. These first two receiver events must be read out of the receiver and discarded. The McBSP will store the digital code until it is read out, so it is important that the McBSP buffer registers be cleared before the DMA is allowed to collect samples. To understand the flowchart and how the hardware behaves, please refer to Figures 5 and 6. The first two cycles in Figure 6 show the configuration word followed by the first channel select command word. The data that is sent out of the ADC during these initial conversion cycles should be trashed because it is from an unknown configuration. The DMA Channel 1 is programmed to synchronize each transfer with an INT3# event. When the ADC de-asserts INT3#, DMA Channel 1 will be prompted to trigger another conversion cycle. When a McBSP1 receive event occurs, DMA Channel 0 reads the receive data register. Figure 5 is a close-up of a single conversion cycle in this single-shot conversion mode. These events will continue until the DMA has completed the number of transfers specified in the element count register for the channel. For example, if 256 samples were desired, the element count register, both for DMA

Figure 4. Single-shot conversion flowchart

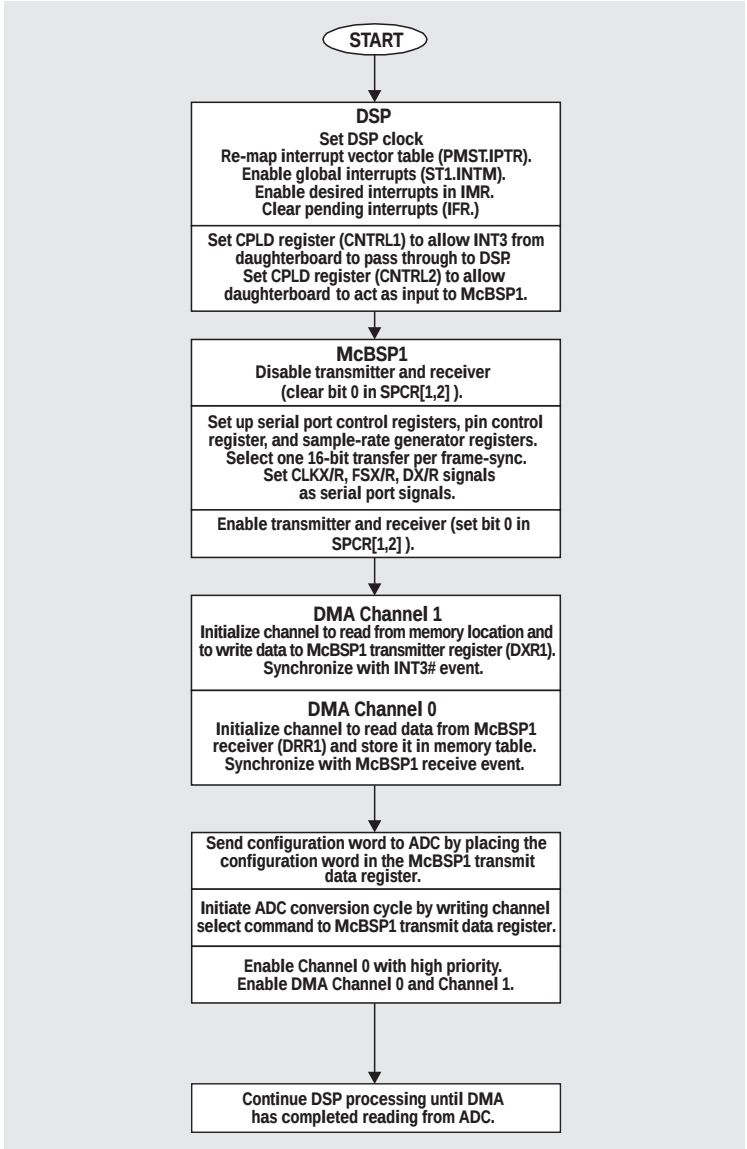

Begin by setting the CPU clock speed. On the C5402 DSK the maximum CPU clock is 100 MHz. If there are any pending interrupts, they are cleared by writing all ones to IFR. Enable DMAC0 interrupt in IMR. The interrupt service routine (ISR) associated with this interrupt can be used to initiate DSP processing of the samples. The interrupt vector table (IVT) needs to be re-mapped so that it will point to the user ISR. The vector table should be placed at the beginning of a data page. The C5402 DSK I/O lines are managed by a CPLD. The CPLD therefore must be initialized to allow the external interrupt 3 (INT3#) signal from the expansion bus to pass through to the DSP. Similarly, the input to the McBSP must be set to arrive from the expansion bus. Before the McBSP registers are configured, it is important that the transmitter and receiver portions be disabled. Once that is accomplished, the user needs to configure the device for the desired operation. Once register values are set, then the transmitter and receiver may be enabled. For this interface the McBSP needs to be configured as in Table 1.

Once all five registers of each DMA channel are configured, the channels are enabled. DMA Channel 0, the channel that stores the received data, is given the highest priority and is synchronized with the McBSP1 receive event. DMA Channel 1, the channel that writes the command word, is synchronized with the INT3# event. The DMA should be programmed as described in Table 2.

The remaining steps are to configure the ADC and trigger the first conversion. This must be done manually by the CPU. The CPU puts the configuration word (for single-channel conversion) in the data transmit register, then places the channel select command on DXR1. Finally, DMA Channel 0 and Channel 1 are enabled. The resulting conversion is communicated to the DSP/DMA with INT3# going low. When INT3# is low, it's time for the DMA to send the next channel select command and read out the received data. The result of this conversion cycle is presented in the next cycle read. The CPU can be off doing its tasks until the DMAC0 interrupt occurs, at which point all the desired samples have been collected, and processing on those samples can begin.

receive and transmit channels, would be initialized to 255. Once the DMA channels have transferred 256 times, the conversion stops. The DMA can be programmed to interrupt the CPU once it is done. Again, the advantage of using the DMA to control the interface to this slow ADC is that the CPU is free to do other tasks.

This interface method and these register settings can be adapted for use with many of TI's current serial data converters that provide an interrupt signal and/or resolutions of ≤16 bits.

References

For more information related to this article, you can download an Acrobat Reader file at www-s.ti.com/sc/techlit/*litnumber* and replace "*litnumber*" with the **TI Lit. #** for the materials listed below.

Document Title **TI Lit. #**

1. Interfacing the TLV2544/TLV2548 ADC to the TMS320C5402 DSP slaa093
2. TLV2548/TLV2544 Datasheet slas198
3. TLVX544/2548EVM Evaluation Module for the TLVX544/TLV2548 10-bit and 12-bit ADC slau029
4. TMS320C54X Assembly Language Tools User's Guide spru102
5. TMS320C54X DSP CPU and Peripherals Reference Set, Vol. 1 . . . . . .spru131
6. TMS320C54XX DSP Enhanced Peripherals Reference Set, Vol. 5 . . . . . .spru302
7. TMS320C54X DSP Reference Set, Vol. 3: Algebraic Instruction spru179
8. TMS320C54X Optimizing C Compiler UG . .spru103
9. Using TMS320C5402 DMA Channels to Read from the TLV2548 . . . . . .slaa095

Related Web sites

www.dataconverter.com
www.ti.com/sc/select
www.ti.com/sc/docs/products/analog/tlv2548.html

Figure 5. Single-shot mode: One conversion cycle

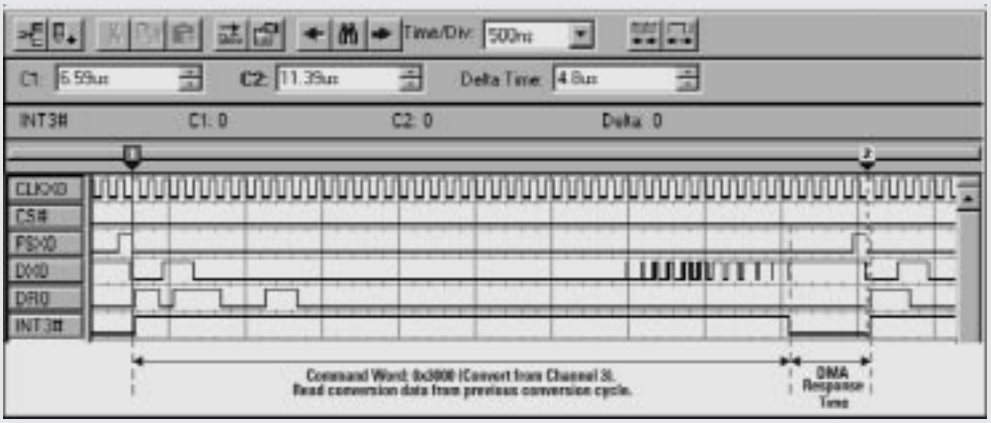

Figure 6. Single-shot conversion: Multiple cycles

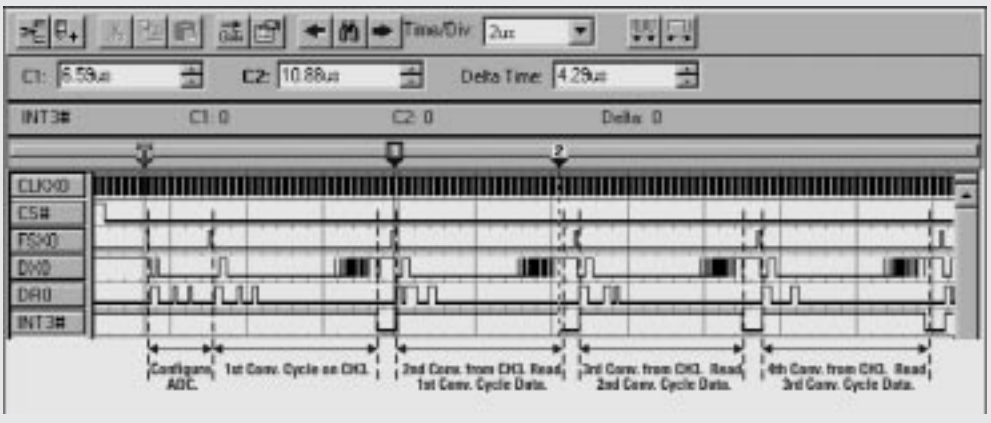

Advantages of using PMOS-type low-dropout linear regulators in battery applications

By Brian M. King
Applications Specialist

Introduction

The proliferation of battery-powered equipment has increased the demand for low-dropout linear regulators (LDOs). LDOs are advantageous in these applications because they offer inexpensive, reliable solutions and require few components or little board area. The circuit model for a typical LDO consists of a pass element, sampling network, voltage reference, error amplifier, and externally connected capacitors at the input and output of the device. Figure 1 shows the circuit blocks of a typical linear regulator. The pass element is arguably the most important part of the LDO in battery applications. The technology used for the pass element can increase the useful life of the battery.

The pass element can be either a bipolar transistor or a MOSFET. The general difference between these is how the pass element is driven. A bipolar pass element is a current-driven device, whereas the MOSFET is a voltage-driven device. In addition, the pass element can be either an N-type (NPN or NMOS) or a P-type (PNP or PMOS) device. N-type devices require a positive drive signal with respect to the output, while P-type devices are driven from a negative signal with respect to the input. Generating a positive drive signal becomes difficult at low input voltages. As a result, LDOs that operate from low input voltages typically are implemented with P-type devices.

When an LDO is selected for a particular application, there are several factors that must be considered. These factors include the dropout voltage (V_{DO}), ground current, noise, input voltage, and thermal response. In low-voltage battery applications, there are two basic factors that make PMOS pass elements much more attractive than PNP pass elements. These two factors are the ground current and dropout voltage. The low ground current and low dropout voltage of PMOS devices can extend the useful life of a battery. Texas Instruments has a wide variety of LDOs containing PMOS pass elements.

Figure 1. Components of a typical linear regulator

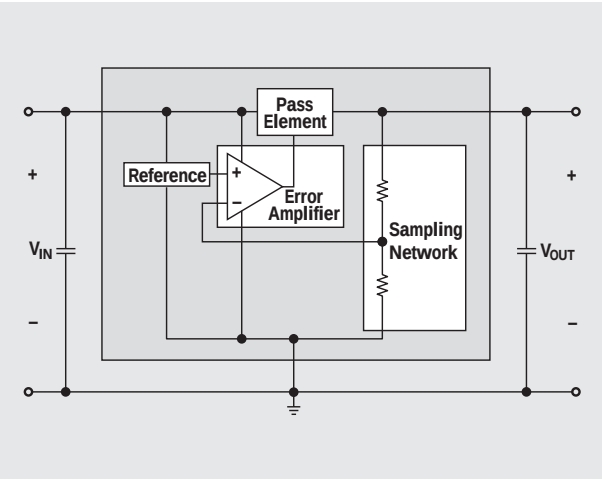

Figure 2. Comparison of typical PNP and PMOS LDO voltages

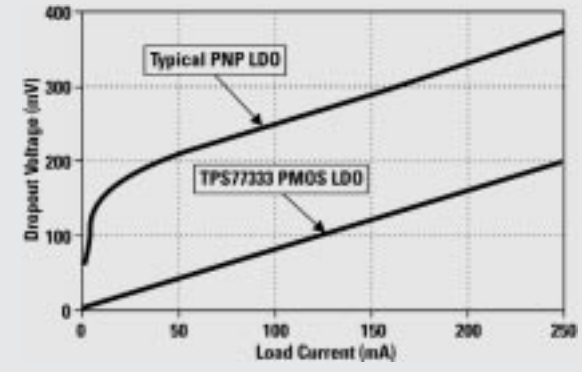

Dropout voltage

The *dropout voltage* is the minimum input-to-output differential voltage required to maintain output voltage regulation. In other words, the minimum allowable input voltage is equal to the sum of the output voltage and the dropout voltage.

The collector-to-emitter voltage drop (V_{CE}) of the transistor determines the dropout voltage of an LDO that contains a PNP transistor as the pass element. The V_{CE} of any given PNP LDO is determined by the physical design of the transistor. In some PNP LDOs, the gain of the transistor is decreased to achieve a lower V_{CE}. However, decreasing the gain results in a higher base current and, consequently, a higher ground current. Thus, in order to achieve a lower dropout voltage in PNP LDOs, a higher ground current usually is compromised. The PNP collector-to-emitter voltage drop increases as the collector current increases. As a result, the dropout voltage of a PNP LDO is dependent on the load current.

As a PMOS LDO encounters a dropout condition, the PMOS pass element becomes linear. When a PMOS device is saturated, the pass element acts as a fixed, low-value resistor. This resistance is referred to as the drain-to-source

on resistance ($R_{ds(on)}$). The $R_{ds(on)}$ of a PMOS device is determined by the physical design of the MOSFET and the available gate drive. The dropout voltage for PMOS LDOs is equal to the output current times the $R_{ds(on)}$ of the MOSFET. Thus, the dropout voltage of PMOS LDOs is also dependent on the output current.

In lower-current applications, PMOS LDOs typically have a lower V_{DO} than that of PNP LDOs. Figure 2 compares the dropout voltage of a PNP LDO with that of a PMOS LDO. The curve for the PNP LDO is typical of a fixed +3.3-V output, 250-mA device. The PMOS curve characterizes the dropout voltage of the Texas Instruments TPS77333, which also produces +3.3 V at 250 mA.

Ground current

Ground current is the difference between the input and output currents and is returned to the input supply through the ground pin of the LDO. The ground current is actually the sum of all the internal bias currents of the LDO. These currents include the reference, sampling network, error amplifier, and pass-element drive currents. Figure 3 shows the locations of these currents. Since the ground current decreases the efficiency of the LDO, it is desirable to minimize this current. The current drawn by the reference, sampling network, and error amplifier is typically quite low, usually less than 100 µA. The drive current for the pass element, however, varies substantially, depending on the type of pass element. In LDOs that contain a PNP transistor, the pass-element drive current makes up a majority of the ground current.

The *quiescent current* of an LDO is defined as the current drawn by the LDO with the output current set to zero. Some vendors, including Texas Instruments, use this term to describe the ground current. When LDO devices are compared, specific attention should be given to the terminology and test conditions used by the different vendors.

Figure 4 shows the curves for a typical PNP transistor in an LDO. The drive current for a PNP transistor is essentially equal to the base current (I_B in Figure 4), and the LDO output current is essentially equal to the collector current (I_C). During normal operation of the LDO, the PNP transistor is operated in the active region shown in Figure 4. In this region, the drive current for a PNP is proportional to the output current by the current gain of the transistor, β. For an average bipolar pass-element transistor, the current gain is between 20 and 100. This means that the drive current is 20 to 100 times smaller than the required output current.

Figure 5 shows the curves for a typical PMOS transistor in an LDO. The LDO output current corresponds to the drain current (I_D). As with a PNP LDO, a PMOS LDO operates in the active region during normal conditions. In this region, varying the gate-to-source voltage (V_{GS}) controls the output current.

The drive current for a PMOS device is determined by the parasitic leakage of the device and is nearly unmeasurable. This results in a much lower ground current than that

Continued on next page

Figure 3. Typical LDO ground current

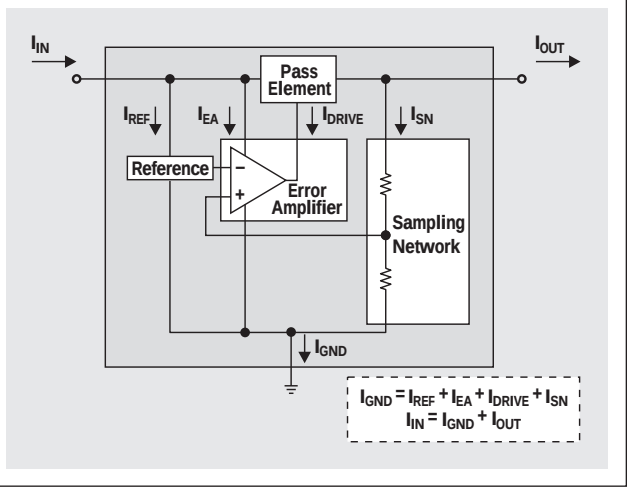

Figure 4. Typical PNP pass-element curves

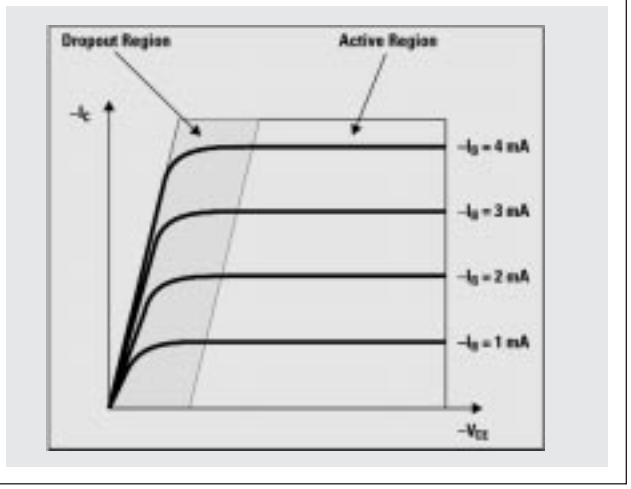

Figure 5. Typical PMOS pass-element curves

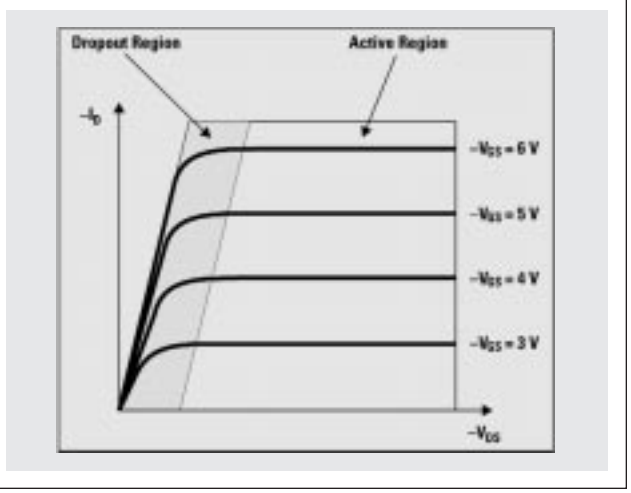

Continued from previous page

of a PNP LDO. Figure 6 compares the ground current of a typical 3.3-V, 250-mA PNP LDO with that of the TPS77333. A 3.75-V input voltage is assumed in this figure. The dependency of the PNP ground current on the load current is evident in this plot, while the TPS77333 ground current remains constant over the entire load range.

In an LDO containing a PNP transistor, the input voltage will affect the ground current. As the voltage drop across the pass element decreases (i.e., as the input voltage approaches the output voltage), the PNP transistor will enter the dropout region shown in Figure 4. When this happens, the transistor requires more base current in order to sustain the output current; i.e., the current gain of the transistor decreases. When operated in a dropout condition, the current gain of a PNP LDO can decrease to a value of less than 1. As a result, the ground current of PNP LDOs increases when a dropout condition is encountered. By comparison, the leakage current of a PMOS device remains fairly constant as the LDO enters the dropout region.

Noise

There are two main characteristics that are of concern when it comes to noise in LDOs. These are the internally generated noise and the power supply ripple rejection (PSRR).

Some designers avoid using a PMOS LDO under the perception that it generates more noise than a PNP LDO. In actuality, the reference and the amplifier are the primary sources of noise, not the pass element. Device manufacturers can reduce the amount of noise generated by an LDO by improving the design of the reference and amplifier.

The PSRR is a measure of how well an LDO attenuates input ripple across a band of frequencies (usually 10 Hz to 1 MHz.) The PSRR is determined by the output impedance of the pass element and the ability of the LDO amplifier to reject supply noise. PNP LDOs typically exhibit better PSRR characteristics than PMOS LDOs. However, the

Figure 6. Comparison of typical PNP and PMOS LDO ground currents

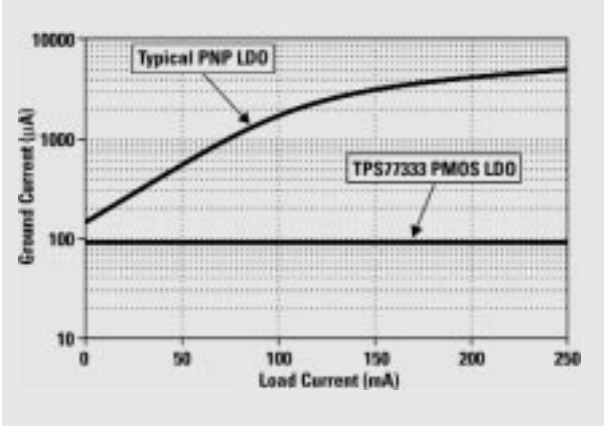

PSRR of an LDO can be improved by increasing the gain of the feedback path or decreasing the ESR of the output capacitor. Selecting an LDO with a well-designed amplifier and proper selection of the output capacitor are the keys to a good PSRR response.

Low input voltage

As mentioned previously, a PMOS pass element is a voltage-driven device. The drive voltage for a PMOS LDO is derived from the input voltage. As a result, the maximum drive voltage is limited by the magnitude of the input voltage. All PMOS devices have a threshold voltage. When the drive voltage drops below the threshold voltage, the PMOS device turns off.

Similarly, even though a PNP transistor is a current-driven device, the emitter-to-base voltage (V_{EB}) of a PNP pass element is derived from the input voltage. In order for a PNP pass element to conduct current, the input voltage must be greater than the V_{EB} of the transistor. However, since the transconductance of a PNP pass element is higher than that of a PMOS pass element, a PNP LDO can be designed to operate at a lower input voltage than a PMOS LDO. In applications where the input voltage may be less than 1.5 V, the use of a bipolar LDO should be considered.

Effect on battery life

As mentioned previously, LDOs that contain PNP transistors as the pass element have a higher ground current than those that contain PMOS transistors. The input current to the LDO is equal to the sum of the ground current and output current. Over the life of a battery, the higher ground current of a PNP LDO will drain the charge more quickly than will a PMOS LDO.

Consider, for example, a 3.3-V application that draws a continuous current of 100 mA from a rechargeable lithium ion battery that has the discharge characteristics shown in Figure 7. If a PNP LDO with the ground current characteristics shown in Figure 6 were used, the ground current would be about 1.7 mA. Thus, the total input current to

Figure 7. Battery-discharge characteristics at 100 mA

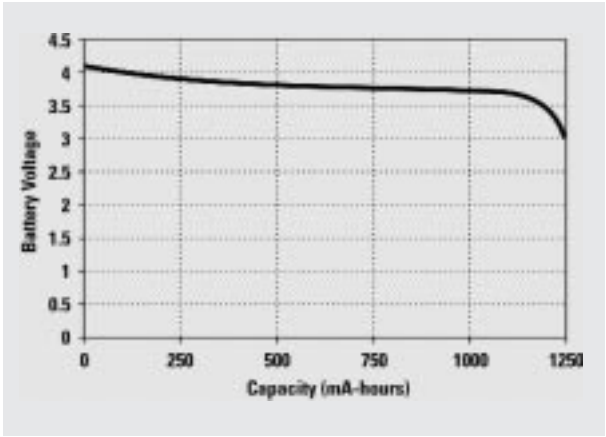

the PNP LDO would be about 101.7 mA. By comparison, if a TPS77333 were used in this application, the ground current would be about 90 µA, and the total input current would be about 100.09 mA. The discharge of the battery versus operating time for these examples is shown in Figure 8. As shown in Figure 8, the larger input current of the PNP LDO will discharge the battery faster than will the PMOS LDO.

The dropout voltage of the LDO affects how much operating time is available from the battery. As can be seen in Figure 2, the PNP LDO has a dropout voltage of about 250 mV at 100 mA, while the TPS77333 has a dropout voltage of around 80 mV at 100 mA. The LDOs encounter a dropout condition when the battery voltage decays to the output voltage (3.3 V) plus the dropout voltage of the LDO. Thus the PNP LDO starts to drop out at 3.55 V, while the PMOS LDO starts to drop out at 3.38 V. The points where the PNP LDO and TPS77333 drop out are shown in Figure 8. Because of the combined effects of the ground current and dropout voltage, the PNP LDO reaches a dropout condition after 11.6 hours, while the TPS77333 drops out after 12.2 hours. Thus, in this example, using a PMOS LDO would increase the operating time of the battery by over 5%.

Figure 8. Battery voltage versus operating time

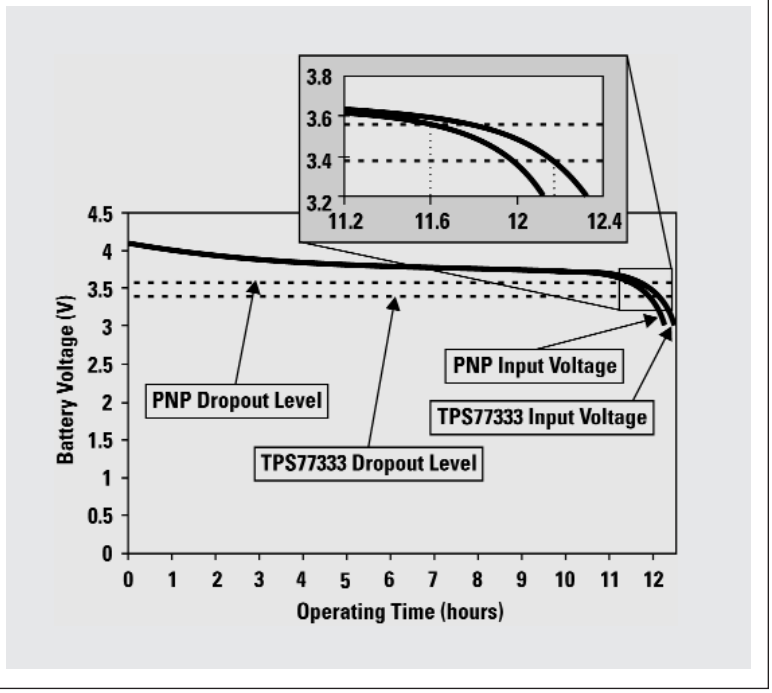

Operation in dropout

PNP-type LDOs can accelerate the decay of battery voltage at the end of battery life. Once the input voltage decreases to a dropout condition, the PNP transistor will saturate, requiring a larger amount of base current. This will increase the ground current and, in turn, increase the load demand on the battery at the end of its life. Consequently, the battery voltage will decrease more sharply after the LDO enters the dropout region. By comparison, since a PMOS transistor is voltage-driven, the drive current remains fairly constant as the battery voltage decays. This effect can become extremely important in applications where the LDO is designed to operate in the dropout region for a majority of the battery life.

Thermal response

Another factor to consider is how the two types of LDOs respond to changes in temperature. The examples given assumed a junction temperature of 25ºC. The dropout voltage of both PNP- and PMOS-type LDOs usually increases as the junction temperature increases. The $R_{ds(on)}$ of a PMOS device is typically 1.7 times higher at 125ºC than at 25ºC. As a result, the dropout voltage of the TPS77333 increases by about 56 mV at 125ºC. The dropout voltage of a PNP LDO at a given load current also increases proportionally to the junction temperature. From the previous example, the V_{DO} of a 3.3-V, 250-mA PNP LDO typically will be 100 mV higher at 125ºC than at 25ºC. At these current levels, PMOS LDOs exhibit a better dropout performance over this temperature range.

The ground currents of PMOS and PNP LDOs behave differently as the junction temperature changes. The current gain of a bipolar transistor typically increases as the junction temperature increases. As a result, the ground current of a PNP LDO may be 50% lower at 125ºC than at 25ºC. For the previous example, it would be reasonable to expect the ground current to be about 900 µA at 125ºC with a 100-mA load. Conversely, the ground current of a PMOS LDO usually remains fairly constant as the junction temperature changes. For example, the ground current of the TPS77333 may rise only from 90 µA to 120 µA at a junction temperature of 125ºC. Even with an increase in junction temperature, the PMOS ground current is an order of magnitude lower than the PNP ground current.

Summary

Because of the nature of the drive current, PMOS LDOs tend to have a significantly lower ground current than PNP LDOs. As a result, PNP LDOs tend to drain the battery voltage more quickly than PMOS LDOs. PMOS LDOs also tend to have a lower dropout voltage than most PNP LDOs. The lower dropout voltage allows the PMOS LDOs to operate from a lower input voltage, which also extends the useful life of the battery. Even over an extended temperature range, the choice of pass-element technology can impact greatly how efficiently a battery is utilized.

Continued on next page

Continued from previous page

Texas Instruments LDO selection guide

TI's LDO product line contains a wide variety of MOSFET, bipolar, and dual LDOs targeted to different output voltages and load requirements. A comparison of the dropout voltages and ground currents of TI's MOSFET LDOs is shown in Table 1. All of these LDOs contain an enable pin that allows the designer to control the startup of the LDO. In addition, some TI MOSFET LDOs contain other features such as a Power Good signal and a RESET (SVS, Power On Reset) signal. The Power Good signal monitors the output voltage and can be used by the circuit designer to generate an alarm or Power OK signal. The RESET signal incorporates a precision delay with the Power Good signal and can be used as a signal in processor applications.

Texas Instruments also offers some dual-output MOSFET LDOs, shown in Table 2. These LDOs are available in a variety of output voltage combinations. In addition, the TPS701 and TPS707 families offer a sequencing function that allows the designer to set the power-up sequence of the two outputs.

Table 3 provides a comparison of the dropout voltages and ground currents of TI's bipolar LDOs. The bipolar LDO family offers output voltages ranging from 1.5 V to 12 V, as well as some adjustable output devices.

Table 1. Texas Instruments MOSFET LDO linear regulators

PART NUMBER	OUTPUT VOLTAGE OPTIONS (V)																MAX. LOAD CURRENT (mA)	TYP. V_{DO} (mV) AT MAX. LOAD	TYP I_q (µA)	AUXILIARY FUNCTIONS			PACKAGING OPTIONS						
	-12	-5	1.224	1.25	1.5	1.6	1.8	2.5	2.7	2.8	3	3.3	3.8	4.85	5	ADJ.				ENABLE	POWER GOOD	RESET	SOT-23	SOIC	MSOP	TSSOP	PowerPAD™	TO-220	TO-263
TPS770xx			•		•		•	•	•	•	•	•			•	•	50	35 to 60	17	•			•						
TPS769xx			•		•		•	•	•	•	•	•			•	•	100	71 to 122	17	•			•						
TPS763xx						•	•	•	•	•	•	•	•		•	•	150	180 to 360	85	•			•						
TPS764xx								•	•	•	•	•					150	300 to 360	85	•			•						
TPS765xx					•		•	•	•	•	•	•			•	•	150	85 to 190	35	•	•			•					
TPS771xx							•	•	•	•	•	•			•	•	150	115 to 150	92	•		•			•				
TPS772xx							•	•	•	•	•	•			•	•	150	115 to 150	92	•	•				•				
TPS74xx					•		•	•			•	•					200	170 to 280	1300	•				•					
UCCx86												•					200	200	10	•				•		•			
UCCx87															•		200	200	10	•				•		•			
UCCx88				•													200	200	10	•				•		•			
TPS72xx								•			•	•		•	•	•	250	190 to 560	180	•	•			•		•			
TPS766xx					•		•	•	•	•	•	•			•	•	250	140 to 310	35	•	•			•					
TPS773xx					•		•	•	•	•		•			•	•	250	200 to 270	92	•		•			•				
TPS774xx					•		•	•	•	•		•			•	•	250	200 to 270	92	•	•				•				
TPS779xx							•	•			•					•	250	250	92	•		•			•				
TPS71025								•									500	330	292	•				•		•			
TPS71Hxx												•		•	•	•	500	146 to 235	285	•	•						•		
TPS71xx												•		•	•	•	500	146 to 235	285	•	•			•		•			
TPS73xx								•			•	•		•	•	•	500	146 to 270	340	•		•		•		•			
TPS775xx					•		•	•				•				•	500	169	85	•		•		•			•		
TPS776xx					•		•	•		•		•				•	500	169 to 285	85	•	•			•			•		
UCCx84-x	•	•														•	500	150 to 200	200	•				•					
TPS777xx					•		•	•				•				•	750	260	85	•		•		•			•		
TPS778xx					•		•	•				•				•	750	260	85	•	•			•			•		
TPS767xx					•		•	•	•	•	•	•			•	•	1000	230 to 500	85	•		•		•			•		
TPS768xx					•		•	•	•	•	•	•			•	•	1000	230 to 500	85	•	•			•			•		
UCCx81-x											•				•	•	1000	500	400	•				•					
TPS751xx					•		•	•				•				•	1500	160	75	•	•						•		
TPS753xx					•		•	•				•				•	1500	160	75	•		•					•		
TPS752xx					•		•	•				•				•	2000	210	75	•		•					•		
TPS754xx					•		•	•				•				•	2000	210	75	•	•						•		
UCCx83-x												•			•	•	3000	400 to 500	400	•								•	•

References

For more information related to this article, you can download an Acrobat Reader file at www-s.ti.com/sc/techlit/*litnumber* and replace "*litnumber*" with the **TI Lit. #** for the materials listed below.

Document Title **TI Lit. #**

1. Tom Kugelstadt, "Fundamental Theory of PMOS Low-Dropout Voltage Regulators," Application Report slva068
2. Bang S. Lee, "Technical Review of Low Dropout Voltage Regulator Operation and Performance," Application Report slva072
3. Bang S. Lee, "Understanding the Terms and Definitions of LDO Voltage Regulators," Application Report slva079
4. James Wong, "A Low-Noise Low Dropout Regulator for Portable Equipment," *Power Conversion and Intelligent Motion* (May 1990), pp. 38-43. —

Related Web sites

http://power.ti.com

www.ti.com/sc/docs/apps/analog/power_supply_support.html

Table 2. Texas Instruments dual-output LDO linear regulators

PART NUMBER	OUTPUT VOLTAGE OPTIONS (V)						MAX. LOAD CURRENT (mA)	TYP. V_{DO} (mV) AT MAX. LOAD	TYP I_q (µA) PER REGULATOR	AUXILIARY FUNCTIONS				PACKAGING OPTIONS	
	3.3 & 1.2	3.3 & 1.5	3.3 & 1.8	3.3 & 2.5	3.3 & ADJ.	ADJ. & ADJ.				ENABLE	POWER GOOD	RESET	SEQUENCING	20-PIN PowerPAD	28-PIN PowerPAD
TPS707xx	•	•	•	•		•	250/125	83	95	•	•	•	•	•	
TPS708xx	•	•	•	•		•	250/125	83	95	•	•	•		•	
TPS701xx	•	•	•	•		•	500/250	170	95	•	•	•	•	•	
TPS702xx	•	•	•	•		•	500/250	170	95	•	•	•		•	
TPS73HDxx			•	•	•		750/750	350	340	•		•			•
TPS767D3xx			•	•	•		1000/1000	350	85	•		•			•

Table 3. Texas Instruments bipolar LDO linear regulators

PART NUMBER	OUTPUT VOLTAGE OPTIONS (V)													MAX. LOAD CURRENT (mA)	TYP. V_{DO} (mV) AT MAX. LOAD	TYP I_q (mA) AT MAX. LOAD	ENABLE	PACKAGING OPTIONS						
	1.5	2.1	2.5	2.8	3	3.2	3.3	3.8	5	8	10	12	ADJ.					SOT-23	SOIC	TSSOP	PowerFLEX™	TO-263	TO-226	TO-220
TPS760xx					•	•	•	•	•					50	120	0.85	•	•						
TPS761xx					•	•	•	•	•					100	170	2.6	•	•						
TL750Lxx									•	•	•	•		150	600	10			•				•	•
TL751Lxx									•	•	•	•		150	600	10	•		•					
TLV2217-33					•									500	500	19				•	•			•
TL750Mxx									•	•	•	•		750	600	60					•			•
TL751Mxx									•	•	•	•		750	600	60	•				•			
TL3317			•	•			•		•				•	3000	1300	5					•			•
UCx82-x	•	•	•										•	3000	350	18						•		•
UCx85-x	•	•	•										•	5000	350	40						•		•

Optimal output filter design for microprocessor or DSP power supply

By Rais Miftakhutdinov
System Engineering, Power Management Products

Introduction

Tight dynamic tolerances for supply voltages of next-generation microprocessors and DSPs at high slew-rate transitions from sleep mode to full-power operation and backwards require fast-transient-response power supplies along with a special decoupling technique. The analysis and optimization of synchronous-buck converters with hysteretic controllers at load-current transients has been presented in References 1 and 2. This article presents a detailed optimization procedure for output filter selection to meet the load-current transient requirements at minimum cost and size. The electrolytic, OS-CON, POSCAP, and ceramic capacitors are compared in a power supply that corresponds to Intel's VRM 8.4 requirements (see Reference 3). These design examples outline the trade-off between cost, size, and efficiency of the power supply and help the user to choose the optimal solution for any particular application.

Figure 1. Analyzed model of power distribution system during load-current transient

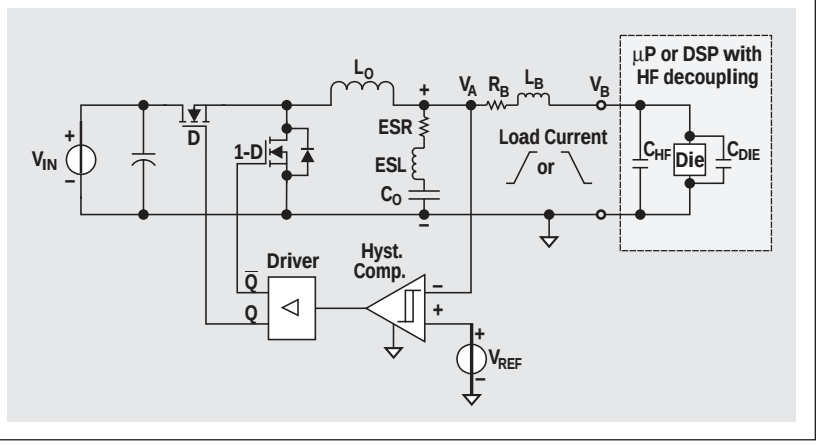

Microprocessor or DSP power supply model

The model shown in Figure 1 presents the microprocessor or DSP power distribution system.

The model includes a synchronous-buck converter with an ideal controller, output inductor (L_O), and output bulk capacitor (C_O) with equivalent series resistance (ESR) and equivalent series inductance (ESL). The equivalent resistor (R_B) represents the added resistance of traces and connectors and characterizes a resistive voltage drop through the supply path. The equivalent inductor (L_B) characterizes an inductive voltage drop through the traces and connectors. The ideal controller has a feedback loop without any delay and limitations on the duty cycle covering the whole possible range from zero to one. In such a case, the converter has minimum peak-to-peak output-voltage transient. The hysteretic controllers from Texas Instruments with relatively small delays and a narrow hysteresis window, such as TPS5210, TPS5211, TPS56XX, TPS56100, and others, are good approximations of the ideal controller with the optimal transient-response characteristics.

Analytical equations were derived for the voltages and currents through the main components of the model in Figure 1 as a function of time both for the load-current step-down and step-up transients. These equations were included in the MATHCAD program to view the voltage and current transient waveforms and to build optimization curves that are described later in this article. To verify the derived equations, the MATHCAD transient waveforms were compared with the measured ones under the same conditions. The measurements were fulfilled on the evaluation board TPS5210SLVP-119. As one can see from Figure 2, the theoretical and measured waveforms are very close for the load-current step-down and step-up conditions.

Impact of system parameters on transients

Dependence on switching-cycle position

The output-voltage transient response depends on the position of the switching cycle when the load-current transient occurs. If the load current steps down, the excessive energy of an output inductor has to be delivered to the output capacitor. The worst case for the step-down transition is when the transient occurs at the end of an upper FET conduction time because the inductor current has its maximum. At this moment the inductor stores the maximum energy while the output ripple voltage also has its maximum. So the transient effect is the most significant at this moment, causing the greatest output voltage spikes in comparison with any other moment (Figure 3).

By contrast, the worst case for the step-up transition is when the transient happens at the end of the switching cycle, because the inductor current and output voltage ripple have their minimum at this moment. Only the output capacitor supplies the load during the step-up transient, while the inductor restores its energy and current to the new load-current level.

Influence of supply-path parasitics

The voltage transient waveforms on the converter output pins (point V_A in Figure 1) and on the microprocessor package supply pins (point V_B in Figure 1) are different because the supply-path resistance (R_B) and inductance

Continued on page 24

Figure 2. Theoretical (a, c) and measured (b, d) waveforms during load-current step-down (a, b) and step-up (c, d) transitions

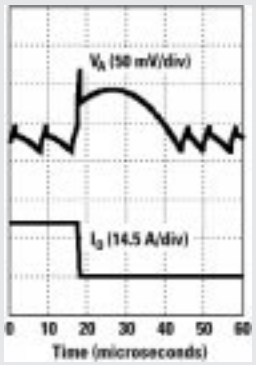

(a) Theory

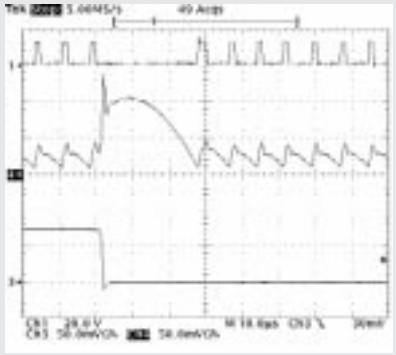

(b) Measurement

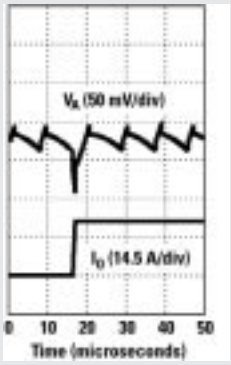

(c) Theory

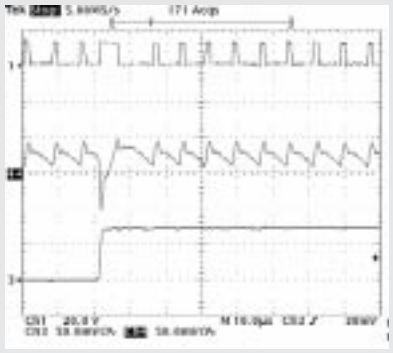

(d) Measurement

Theoretical waveforms show the output voltage (top) and load-current (bottom) transients. Measured waveforms include voltage (V_{DS}) of the low-side FET (Ch1: 20 V/div), output voltage (Ch4: 50 mV/div), and load current (Ch3: 14.5 A/div).

Figure 3. Output voltage (bottom curve) and inductor current (dashed) waveforms for the different instants when the load-current (top, solid) step-down transition occurs

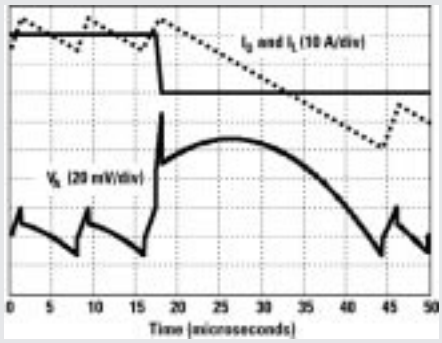

(a) Worst case: Transient occurs at the end of the upper FET's conduction time

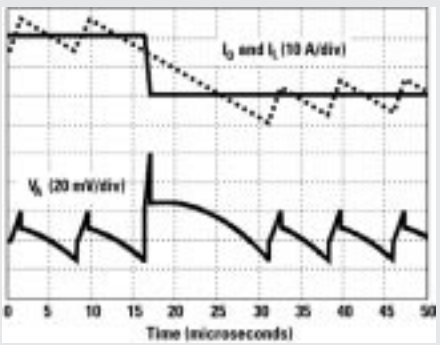

(b) Best case: Transient occurs at the end of the switching cycle

Continued from page 22

(L_B) cause the additional voltage drop. If the output current step (ΔI_O) and slew rate (SR) are defined as

$$\Delta I_O = I_{O(MAX)} - I_{O(MIN)} \text{ and} \quad (1)$$

$$SR = \Delta I_O / t_O , \quad (2)$$

where t_O is the output current transition duration, then the additional voltage drop (V_B) through the supply paths is defined by

$$V_B = VR_B + VL_B = \Delta I_O \times R_B + SR \times L_B . \quad (3)$$

Assume that R_B = 1.5 mohm; L_B = 1.0 nH; ΔI_O = 23.8 A; and SR = 20 A/µs in accordance with the VRM 8.4 requirements. Then the voltage drop through the supply path is

$$V_B = 35.7\,mV + 20\,mV = 55.7\,mV.$$

For the 1.65-V output power supply, this means almost 3.8%. This example shows why it is important to keep the output filter capacitors as close as possible to the microprocessor package to avoid a significant voltage drop due to supply-path parasitics.

Optimal output inductor

It seems obvious that the lower output inductor value enables better transient-response characteristics because of faster inductor current change to the new level after the load-current transient occurs. The example in Figure 4 shows that, in reality, after some optimal point (Figure 4b), further decreasing of the inductor value increases the peak-to-peak transient amplitude because the output ripple rises significantly. As shown later, the optimal inductor value depends on switching frequency and the type of output bulk capacitors.

Two extreme values of an output-voltage transient

Typical load-current transient waveforms are shown in Figure 5. The output-voltage waveform has two extreme values, Vm1 and Vm2. For most applications, the transient slew rate of the load current is much higher than the maximum slew rate of the output inductor current. Because of that, the first extreme value, Vm1, depends mainly on the output capacitor and supply-path parasitics. It is not affected significantly by the controller transient-response characteristics.

The second extreme value, Vm2, depends on resistive components ESR and R_B, capacitive component C_O, inductor value L_O, and the converter characteristics, including switching frequency and type of control. Vm2 does not exist if the following inequality is fulfilled:

$$ESR \times C_O > m \times t_s \times \left(\frac{1}{2} + \frac{\Delta I_O}{\Delta I_L} \right), \quad (4)$$

where t_s is a switching cycle, and ΔI_L is a peak-to-peak ripple portion of the output inductor current. The parameter m depends on the type of transient. For the worst-case step-down transient, m = 1 – D, and for the worst-case step-up transient, m = D. Of course, only the first spike has to be considered in this situation during the design.

Assume that the output filter capacitors are connected in parallel and that each capacitor has the characteristics C_O1, ESR1, and ESL1. The number "1" after a parameter means that that parameter relates to one of many capacitors connected in parallel. It is shown in Reference 2 that if ΔVreq is the maximum allowable peak-to-peak transient tolerance, then the required number of output bulk capacitors, N1 and N2, to meet the conditions Vm1 = ΔVreq and Vm2 = ΔVreq, respectively, can be defined as in Equations 5 and 6 at the bottom of this page. Equations 5 and 6 can be used for the optimal output filter design.

Active droop compensation

One can see from Equations 5 and 6 that the number of capacitors can be lowered by increasing ΔVreq. The active droop compensation is an effective technique to do that. The droop compensation means that the dc output-voltage level of the converter is set to the highest level within the specification window at no-load condition and to the lowest level at full-load. This approach degrades the static load regulation but increases the output-voltage dynamic tolerance by as much as twofold, thus reducing the number of bulk capacitors required. For the same output filter, this technique allows a decrease in the peak-to-peak output-voltage transient response. The popularity of this idea is confirmed by the fact that it has numerous names like "Programmable Active Droop™," "Active Voltage Positioning," "Adaptive Voltage Positioning," "Summing-Mode Control," etc. The transient waveforms with and without active droop compensation are shown in Figure 6. One can see that without droop compensation (Figure 6a), the output-voltage peak-to-peak amplitude is 146 mV and exceeds the requirements, as shown by the cursors. With droop compensation (Figure 6b), the peak-to-peak transient is only 78 mV, keeping the

$$N1 = \frac{\frac{ESL1}{t_O} + ESR1 + \frac{t_O}{2 \times C_O 1} + \left(ESR1 + \frac{t_O}{2 \times C_O 1} \right) \times \left(1 - \frac{t_O}{m \times t_s} \right) \times KL}{\frac{\Delta Vreq}{\Delta I_O} - \frac{L_B}{t_O} - R_B} \quad \text{and} \quad (5)$$

$$N2 = \frac{\frac{1}{2} \times \left[\frac{m \times t_s}{C_O 1} - \frac{t_O}{C_O 1} + \left(ESR1 + \frac{ESR1^2 \times C_O 1}{m \times t_s} + \frac{m \times t_s}{4 \times C_O 1} \right) \times KL + \frac{m \times t_s}{C_O 1} \times \frac{1}{KL} \right]}{\frac{\Delta Vreq}{\Delta I_O} - R_B}, \quad (6)$$

where $KL = \frac{V_{OUT} \times (1-D) \times t_s}{L_O \times \Delta I_O}$, and $D = V_{OUT}/V_{IN}$ is a duty cycle.

Continued on page 26

Figure 4. Transient waveforms with different inductor values (L_O)

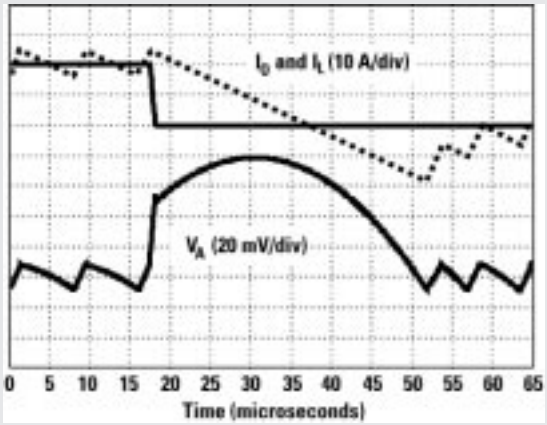

(a) L_O = 1.6 µH, $V_{O(MAX)}$ = 79 mV, t_{recov} = 34 µS

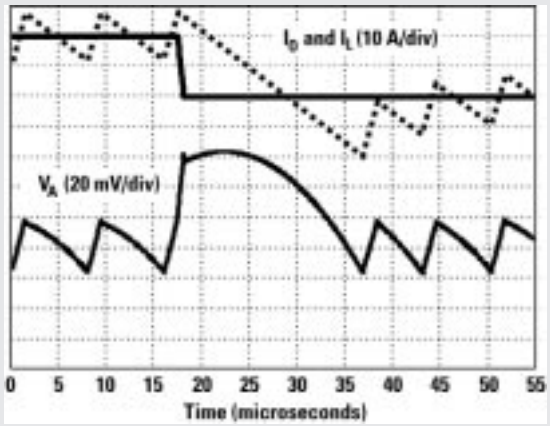

(b) L_O = 0.8 µH, $V_{O(MAX)}$ = 62 mV, t_{recov} = 19 µS

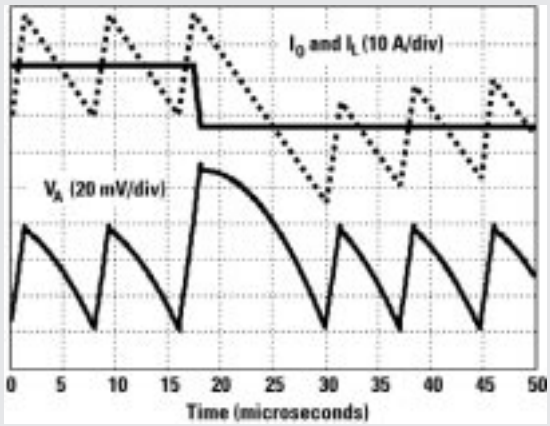

(c) L_O = 0.4 µH, $V_{O(MAX)}$ = 72 mV, t_{recov} = 12.5 µS

Figure 5. Typical load-current transient waveforms

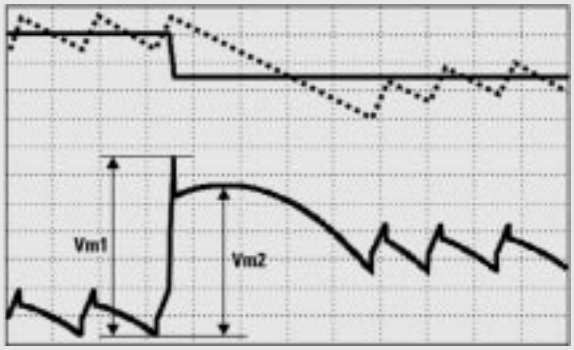

Figure 6. Active droop compensation technique

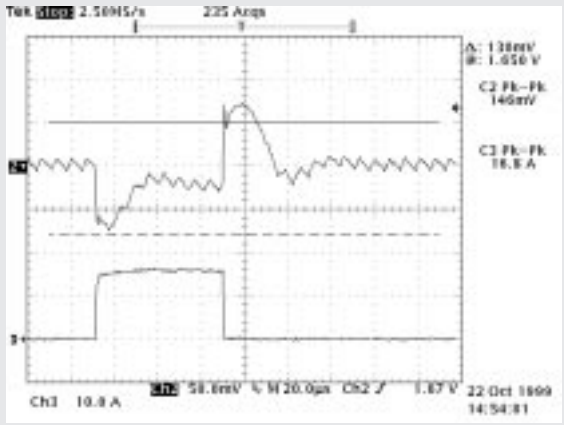

(a) Without droop compensation ($V_{OUT(P-P)}$ = 146 mV)

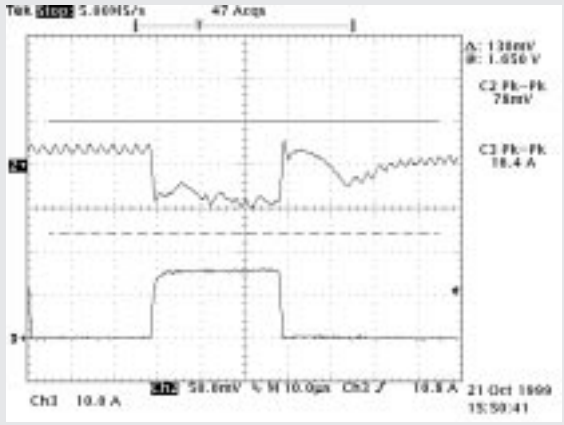

(b) With droop compensation ($V_{OUT(P-P)}$ = 78 mV)

Channel 2 shows output voltage (50 mV/div.), Channel 3 shows load current (10 A/div.), and the cursors show the required limits for the output voltage.

Continued from page 24

output voltage well within the requirements for the same load-current transient conditions.

Optimal output filter selection procedure

The following step-by-step design procedure shows how to select output capacitors, an inductor value, and a switching frequency that are optimized for a specific application. For this design the typical VRM 8.4 requirements are used as an example:

V_{IN} = 5 V, V_{OUT} = 1.65 V, ΔV_{OUT} (dc) = –80 mV ÷ +40 mV, ΔV_{OUT} (ac) = –130 mV ÷ +80 mV, $I_{O(MAX)}$ = 26 A, $I_{O(MIN)}$ = 2.2 A, ΔI_O = 23.8 A, ΔVreq = 180 mV, SR = 20 A/µs, R_B = 1.5 mΩ, and L_B = 1 nH.

1. Definition of the worst-case transient

Select which type of transient, a load-current step-up or step-down, is the most important to optimize. The transient, caused by the load-current transition, is completed when the inductor current has reached the new steady-state current level. The inductor current slew rate depends on the voltage applied to the inductor. This voltage is equal to $V_{IN} – V_{OUT}$ during a load-current step-up, or to V_{OUT} during a load-current step-down. For most microprocessor and DSP applications, usually $(V_{IN} – V_{OUT}) > V_{OUT}$. This means that the worst case is defined by the load-current step-down transition because the lower voltage, V_{OUT}, changes the inductor current more slowly. In such a case the load-current step-down has to be optimized first; then, after the output filter selection, the load-current step-up transient has to be verified to meet requirements.

2. Maximum peak-to-peak dynamic tolerance

An accurate output-voltage budget needs to be done to determine a maximum dynamic output-voltage tolerance, ΔVreq. The dynamic and static supply-voltage limits have to be compared with all potential tolerances, including set-point accuracy, time and temperature variation, and line and load regulation. Use the droop compensation and adjust the nominal output voltage to get the maximum possible ΔVreq. Figure 7 shows the output-voltage budget calculation for this particular example. For the step-down transient, the required window is:

Vreq = 1,730 mV – 1,570 mV – 2 mV – 2 × 6 mV – 50 mV = 96 mV

For the step-up transient, it is:

Vreq = 1,690 mV – 1,520 mV – 2 mV – 2 × 6 mV – 50 mV = 106 mV

The required droop compensation is:

Droop = 106 mV – 16 mV – (1,570 mV – 1,520 mV – 2 mV) = 42 mV

Table 1. Comparison of different types of capacitors

TYPE	VENDOR	PART NUMBER	V_{dc} (V)	CAPACITANCE (µF)	ESR (mohm)	ESL (nH)	SIZE (mm)	RELATIVE COST
Aluminum electrolytic	Rubycon	6.3ZA1000	6.3	1000	24	4.8	∅10 x 16	1
OS-CON	Sanyo	4SP820M	4	820	8	4.8	∅10 x 10.5	6
POSCAP	Sanyo	4TPC150M	4	150	40	3.2	7.3 x 4.3 x 1.9	3
Ceramic	Murata	GRM235Y5V226Z10	10	22	20	0.5	3.2 x 2.5 x 1.35	0.7

Figure 7. Output-voltage budget for VRM 8.4 power supply (not scaled)

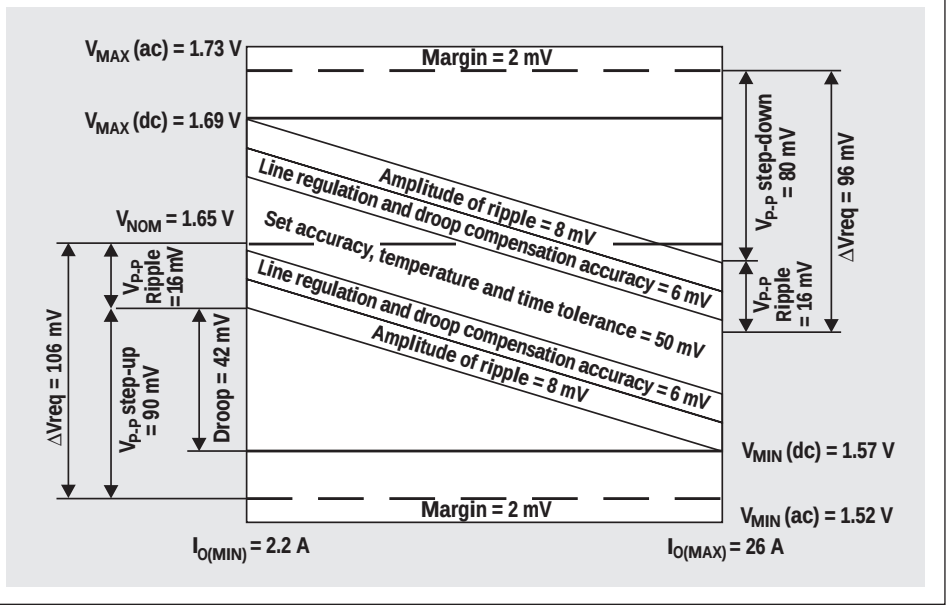

Figure 8. Optimization curves N1 (dashed) and N2 (solid) as a function of output inductance (L_O) and switching frequency (f_s) for different types of capacitors

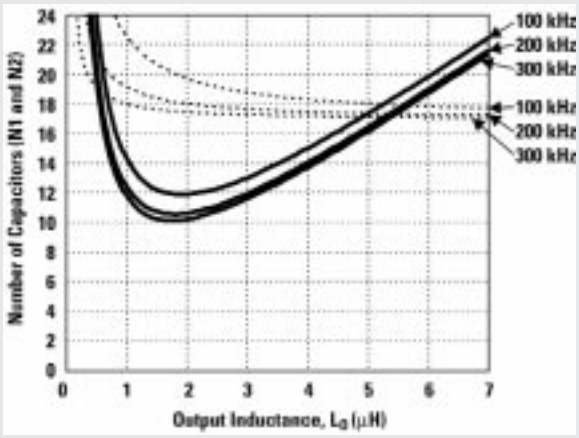

(a) Aluminum electrolytic

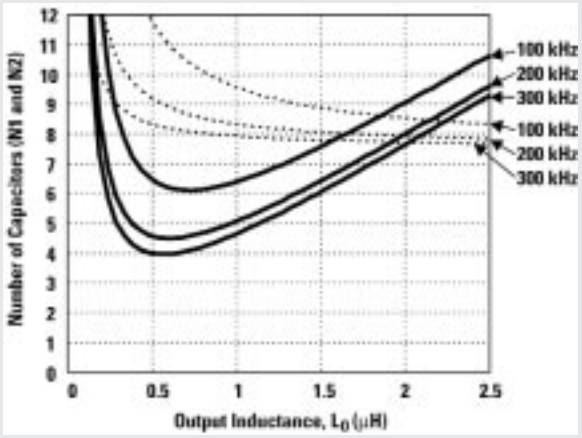

(b) OS-CON

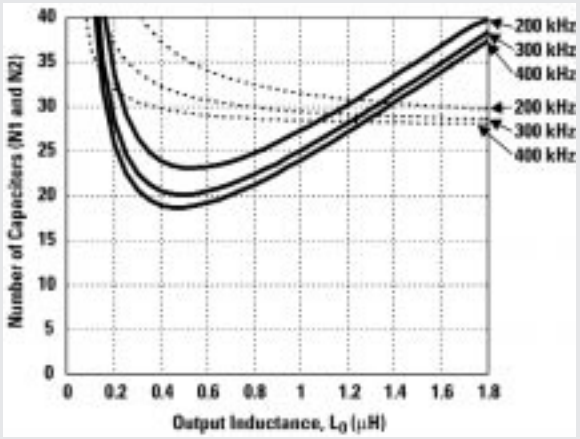

(c) POSCAP

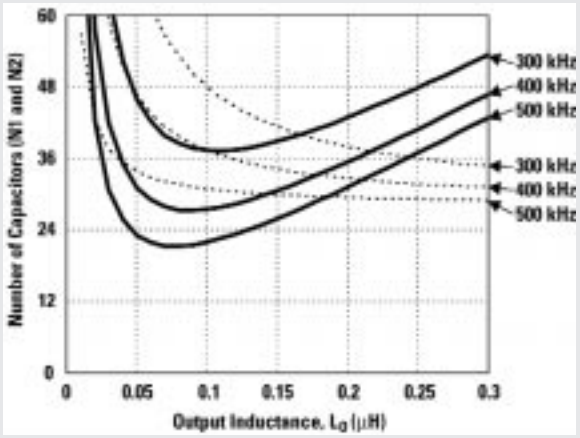

(d) Ceramic

3. Output bulk capacitor selection

Equations 5 and 6 show that the number of output bulk capacitors, N1 and N2, can be decreased if capacitors with low ESL1, ESR1, and high enough capacitance, C_O1, are used. Electrolytic, OS-CON, POSCAP, and ceramic capacitors are the most popular candidates for this application. Table 1 shows the main characteristics of capacitors that have been selected for the comparison in this design. This table does not restrict the list of capacitors and vendors, and the selected capacitors illustrate only the trade-off between different types based on cost, size, reliability, and efficiency. The capacitor vendors usually provide the impedance and ESR curves based on measurements with sinusoidal waveforms. The ESL value usually is not specified. For better design accuracy, the ESR and ESL have to be estimated by measuring capacitor reaction on the high-slew-rate linear charge or discharge current if the capacitor is intended for use in microprocessor or DSP power supplies. Because of this, some numbers in the table may differ from the specification data.

For the ESL and ESR estimate, the parasitic inductance and resistance of the traces and vias required for capacitor mounting have to be included. Usually, many smaller capacitors connected in parallel yield lower ESL and ESR compared with larger capacitors. Small surface-mount capacitors can be located as close as possible to the microprocessor or DSP package. This is important to minimize R_B and L_B values.

4. Output filter optimization curves

The number of capacitors, N1 and N2, as a function of output inductance (L_O) and switching frequency ($f_s = 1/t_s$) are shown in Figure 8 for the electrolytic, OS-CON, POSCAP, and ceramic capacitors.

Continued on next page

Continued from previous page

One can see that the first spike curves, N1 (dashed), do not have an extreme value and increase rapidly at low output inductance. The second spike curves, N2 (solid), have minimum value at some output inductance. In most cases the curves N1 cross the curves N2 at two points. The lowest-integer number of output capacitors, which is still higher than the cross-section of curves N1 or N2, is the minimum number that satisfies the requirements. The inductance has to be selected as close as possible to the cross-section points of both curves, or in-between. To avoid variation in the transient response due to component tolerances, it is wise to select the inductance in the region where the slew rate of the curves is not too high.

It is shown that the number of aluminum electrolytic and OS-CON capacitors does not drop significantly at frequencies higher than 200 kHz. The POSCAP capacitors work effectively at up to 350 kHz. The ceramic capacitors work well at 500 kHz and at higher frequencies. Here, the switching frequency is rather restricted by the power losses in semiconductors.

One can see that, for the aluminum electrolytic and OS-CON and partly for the POSCAP capacitors, the curve N1 is typically higher than the curve N2. This is because they have relatively high ESL1 and large capacitance C_O1. The number of capacitors N1 also rises rapidly if the supply-path stray inductance is too high. The additional high-frequency decoupling helps to decrease equivalent ESL and to reduce the number and cost of bulk capacitors.

In accordance with this design example, aluminum electrolytic and OS-CON capacitors in the 100- to 200-kHz switching-frequency range and with the output inductance value around 1 to 2.0 µH are preferable for applications requiring low power losses. The POSCAP capacitors have lower ESL1, but their number is higher because of relatively high ESR1 and low capacitance C_O1. Their preferable application is low-height DC-DC converters with a switching-frequency range from 250 to 300 kHz. The optimal value for the output inductor is around 0.7 µH. The number of ceramic capacitors might be too large at frequencies lower than 500 kHz. Their preferable application area is minimum-size, high-frequency converters. One can see that their ESL1 is very low, but the impact of the second extreme value is significant because of the low capacitance C_O1. The system cost, temperature range, available space, reliability, cooling conditions, and life of the product have to be considered during final selection of the output filter.

Assume that the aluminum electrolytic capacitor has been selected for further consideration. Figure 8a shows that, at 200-kHz switching frequency and with a 2-µH inductor, the required number of capacitors is 18 because of the high first spike. The number of electrolytic capacitors in this case can be decreased if a few high-frequency decoupling capacitors are added to decrease the impact of ESL and L_B. Seven 805-size ceramic capacitors of 1 µF each have been added. Each capacitor has an ESL1 of 2.6 nH, including inductance of vias and traces. The equivalent inductance of 7 capacitors placed very close to the microprocessor is 2.6 nH/7 = 0.37 nH. For this design, the load-current slew rate and supply-bus inductance are SR = 20 A/µs and

Figure 9. Optimization with aluminum electrolytic capacitors

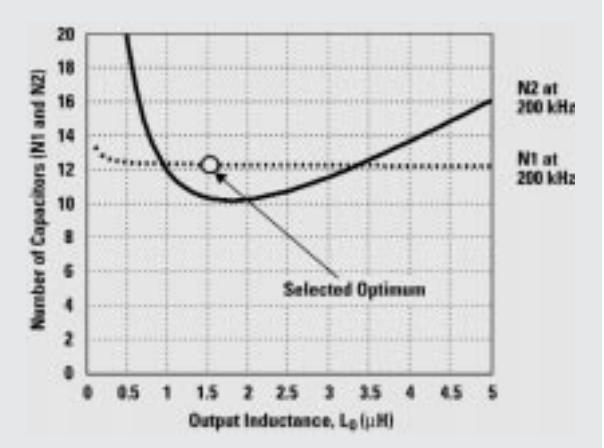

New optimization curves N1 and N2 for the aluminum electrolytic capacitors with additional high-frequency decoupling

L_B = 1 nH. Adding the high-frequency decoupling capacitors decreases the slew rate roughly three times in this case:

$$SR_{new} = SR \times (0.37\ nH/1\ nH) = 7.4\ A/\mu s$$

The new optimization curves for the aluminum electrolytic capacitors with high-frequency decoupling are shown in Figure 9. One can see that curves N1 and N2 are much closer to each other because the high-frequency decoupling lowers the effect of inductive parasitics. The optimal inductor value, 1.5 µH, is selected, and the number of capacitors is 12 instead of 18 without the additional 7 ceramic capacitors of 1 µF each.

The transient waveforms based on this design are shown in Figure 10. One can see that the output voltage at low load is shifted at higher levels because of active droop compensation. Both step-up and step-down transients are acceptable for microprocessors in accordance with the VRM 8.4 requirements.

Conclusion

A power-supply system for powering high-slew-rate transient loads, such as a microprocessor or DSP, was analyzed. The selected model, based on practicality and accuracy, included a synchronous-buck converter with the controller, output inductor, output bulk capacitor with parasitics, and power-supply traces between the bulk capacitors and the microprocessor or DSP package. The accuracy of the model and derived equations was confirmed by comparison with the measurement results. It was shown how the different parameters of the model influence transient-response characteristics. A step-by-step optimal design procedure for the minimum-size and least costly output filter was suggested based on the derived equations. A design example of a DC-DC converter in accordance with the VRM 8.4 requirements was presented, and different types of bulk capacitors like aluminum electrolytic, OS-CON, POSCAP, and ceramic were compared.

Figure 10. Transient waveforms of an optimized output filter with aluminum electrolytic capacitors

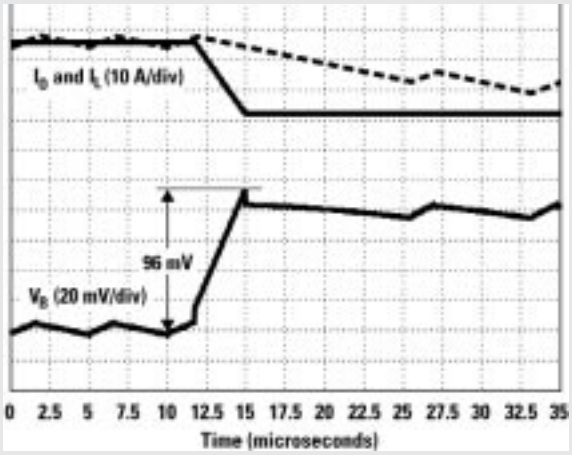

(a) Load-current step-down

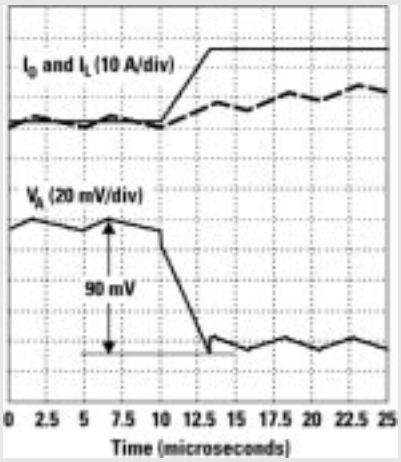

(b) Load-current step-up

References

1. R. Miftakhutdinov, "Analysis of Synchronous Buck Converter with Hysteretic Controller at High Slew-Rate Load Current Transients," *Proc. of High Frequency Power Conversion Conference* (1999), pp. 55-69.
2. R. Miftakhutdinov, "Analysis and Optimization of Synchronous Buck Converter at High Slew-Rate Load Current Transients," *Proc. of Power Electronics Specialists Conference* (2000), pp. 714-720.
3. Intel Corporation, "VRM 8.4 DC-DC Converter Design Guidelines," November 1999, order number 245335-001.

Related Web sites

http://power.ti.com
www.ti.com/sc/docs/tools/analog
www.ti.com/sc/select

Performance of LVDS with different cables

By Frank Dehmelt
Application Engineer, Data Transmission

Introduction to LVDS

Low-voltage differential signaling (LVDS) runs fast—very fast. One of the most frequently asked questions about data transmission applications is, "How fast and how far?" The answer depends on technology, system circumstances (noise, crosstalk, stubs, etc.) and the connection media.

TI's family of LVDS line circuits enables signaling rates in excess of 400 Mbps, and using such TIA/EIA-644 standard devices may result in cable performance being the determining factor in the overall system performance. LVDS is a data transmission standard that utilizes a balanced interface and a low voltage swing to solve many of the problems associated with existing signaling technologies. Lower signal amplitudes reduce the power used by the line circuits, and balanced signaling reduces noise coupling to allow higher signaling rates. LVDS, as standardized in TIA/EIA-644, specifies a maximum signaling rate of 655 Mbps. In practice, the maximum signaling rate will be determined by the quality of the transmission media between the line driver and receiver. Since a transmission line's length and characteristics determine the maximum usable signaling rate, this article looks at some of the dependencies and interactions between these cable characteristics and the signaling rate. See Reference 1 for the detailed version of this article.

Cable selection

Prior to the cable selection, a designer has to evaluate the determining system parameters such as:

- signaling rate,
- cable length,
- single-ended or differential (balanced) signaling,
- point-to-point, multidrop or multipoint configuration,
- noise margin,
- flexibility, and
- costs.

Depending on the specific application and environment, the following decisions need to be made:

- Unshielded or shielded (taped, braided, or combination of both)?
- Round or flat?
- Coaxial, multiconductor or twisted pair (TP) cable?

The need for shielding depends mainly on the noise environment. For long transmission lines, a braided or served shield is recommended to ensure good isolation between the signal lines and the environment. However, this type of shielding is permeable at high frequencies, and double-shielded cables that are both taped and braided typically perform better. Multiconductor cables are cheaper and easier to handle than twisted pair or coaxial cables, especially in terms of termination. While twisted pair is less expensive and more flexible than coaxial, it generally does not provide the noise immunity and bandwidth available with coaxial cables. Nevertheless, differential data transmission requires a balanced pair of conductors. The answer to the round- or flat-cable question is usually determined by the environment. For internal applications with low noise, a flat untwisted cable is usually adequate. However, in noisy environments, shielding is often required, and industry standards call for shielded twisted pair (STP) cable. For balanced (or differential) data transmission, such as LVDS, twisted pair cable is recommended since it provides two identical conductors to transmit the signal and its complement. Ideally, any distortion will affect both conductors equally; therefore, the differential signal will not change.

The cable-standard TIA/EIA-568-A

Since cable quality contributes strongly to signal quality, it should be evaluated in detail. One standard, the TIA/EIA-568-A Commercial Building Telecommunications Cabling Standard, defines the transmission requirements for commercial building telecommunication wiring. Twisted pair is classified herein in different categories, abbreviated by CATX. CAT3 is characterized up to 16 MHz, CAT4 to a maximum of 20 MHz, and CAT5 for 100 MHz and above. CAT6 and CAT7 are in preparation. Parameters such as attenuation, dc resistance, skew, capacitance to GND and between lines, etc., are specified in TIA/EIA-568A.

Measurements

Seven different cables are tested with the LVDS evaluation module (EVM). Each EVM contains one SN65LVDS31 quad line driver and one SN65LVDS32 quad line receiver, and each of the cables listed below is tested as the interconnection media between the LVDS driver and receiver.

- Cable A: CAT 3, no shield, outside conductor diameter ∅ 0.52 mm
- Cable B: CAT 5, no shield, ∅ 0.52 mm
- Cable C: CAT 5, taped over all shield, ∅ 0.52 mm
- Cable D: Exceeding CAT 5, specified up to 300 MHz, braided over all shield plus taped individual shield for any pair, ∅ 0.64 mm
- Cable E: Exceeding CAT 5, specified up to 350 MHz, ∅ 0.64 mm, no shield
- Cable F: Exceeding CAT 5, specified up to 350 MHz, self-shielded, ∅ 0.64 mm
- Cable G: Twin-axial cable, specified up 1 GHz

For each measurement, a pseudo-random binary signal (PRBS) with a non-return to zero (NRZ) format is used. PRBS patterns are applied to the input of the transmitter; then eye patterns are measured at the input of the receiver. Tests are performed on Cables A through F with lengths of

1, 5, and 10 meters (only a 10-meter length is available for the testing of Cable G). Since all cables tested contain four pairs, crosstalk is created by transmitting through two of the pairs in one direction while the remaining two pairs are driven in the opposite direction. All of the data listed is measured with the four transmitters in operation.

Test setup

The LVDS EVM contains one SN65LVDS31 quad line driver and one SN65LVDS32 quad line receiver, as shown in Figure 1. All four channels of the line driver are utilized to simulate crosstalk by transmitting PRBS in opposite directions. The EVM is CE-certified and available via distribution, and detailed information on the EVM is available in Reference 2.

Jitter measurement

The eye pattern is a useful tool to measure the overall signal quality at the end of a transmission line. It includes all of the effects of systemic and random distortion, and shows the time during which the signal may be considered valid. A typical eye pattern is illustrated in Figure 2 with the significant attributes identified.

Several characteristics of the eye pattern indicate the signal quality of the transmission circuit. The height or opening of the eye above or below the receiver threshold level at the sampling instant is the noise margin of the system. The spread of the transitions across the receiver thresholds measures the peak-to-peak jitter of the data signal. The signal rise and fall times can be measured relative to the 0% and 100% levels provided by the long series of low and high levels.

Jitter is the time frame during which the logic state transition of a signal occurs. The jitter may be given either as an absolute number

Continued on next page

Figure 1. Test setup

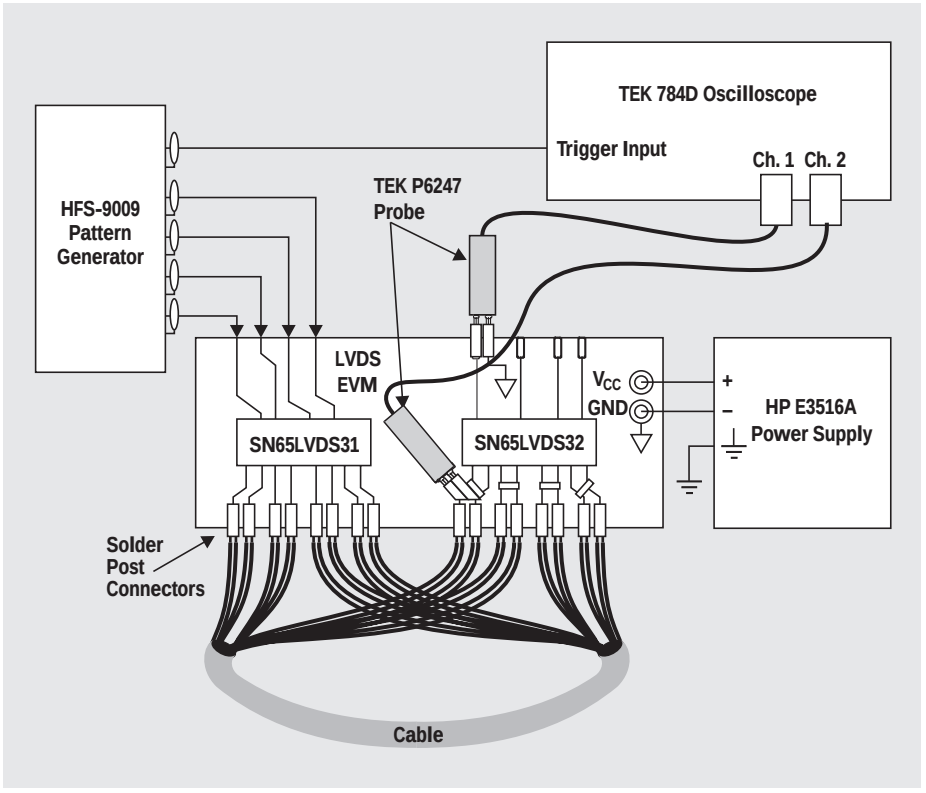

Figure 2. Typical eye pattern

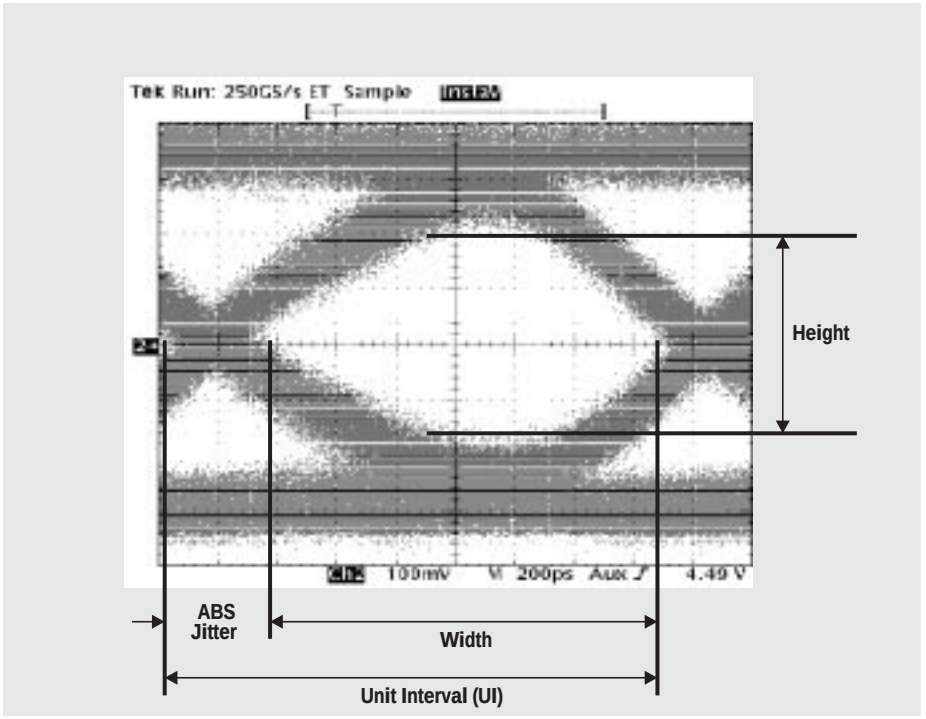

Continued from previous page

or as a percentage with reference to the unit interval (UI). This UI or bit length equals the reciprocal value of the signaling rate, and the time during which a logic state is valid is just the UI minus the jitter. Percent jitter (the jitter time divided by the UI times 100) is more commonly used, and it represents the portion of UI during which a logic state should be considered indeterminate.

Table 1. Signaling rates vs. cable length for 5% jitter

CABLE LENGTH (m)	SIGNALING RATES (Mbps)						
	CABLE A	CABLE B	CABLE C	CABLE D	CABLE E	CABLE F	CABLE G
1	240	200	240	270	180	230	N/A
5	205	210	230	250	215	230	N/A
10	180	150	195	200	145	180	195

Results

Jitter at the input of the receiver is measured at the zero voltage differential, then calculated with respect to the duration of the unit interval. The results are expressed as a percentage of jitter. At 400 Mbps the jitter ranges between 17% for the worst twisted pair cable (mostly 12% to 13%) and 10% for the twin-axial cable over 10 meters. For shorter runs the jitter is reduced towards 8% to 13% for 5 m of length and ranges from 7.5% to 12.5% over distances of 1 m. The decrease of jitter with a reduced data rate is a linear function, so half the data rate equals approximately half the amount of jitter. Detailed graphs on these measurements can be found in Reference 1.

The linear increase of the jitter (as a percent of UI) with the signaling rate in all measured cables is a relative measure of the high-frequency characteristics of the cable. The results are summarized in Table 1, which displays the signaling rates that resulted in a jitter of 5% of UI present at the input to the LVDS32 receiver. System tolerance to jitter is highly application-dependent, and maximum allowable jitter tolerances typically range from 5% to 20% depending upon actual system requirements. Note that this data was collected with signals present on the other three twisted wire pairs in the cable.

Test results show that, as expected, slightly better performance was achieved with the shielded Cables E and F than the unshielded Cables C and D. It is difficult to identify the noise coupling source as inter-system or intra-system, but if electromagnetic noise is a concern, shielding should be used. Transmitting data through a single channel and signal pair may reduce the absolute jitter by up to 10%.

Cable length and signaling rate

Equation 1 was developed based on the gathered data. This equation approximates percent output jitter through an LVDS32 receiver, given the cable length (m) and signaling rate (Mbps), and is valid for cable lengths from 5 m to 20 m and signaling rates of 100 Mbps up to 400 Mbps.

$$\text{Output jitter} = \frac{\left\{\left[1+\left(\frac{0.0023(S)}{1+\frac{L-1}{7.63}}\right)\right][(200+15L)S]\right\}}{10{,}000}, \quad \textbf{(1)}$$

where S is the signaling rate in Mbps and L is cable length in meters.

This data is based upon the CAT5 cables tested, and the reader should be advised that a marginal cable that minimally meets the requirements of CAT5 may yield actual performance less desirable than the results predicted here with Equation 1.

Conclusion

Especially at high data rates like with LVDS, the quality and length of the transmission media have a significant impact on the overall system performance. For signaling rates in the range of several hundred megabits, one has to expect a jitter ranging around 10% to 15% on distances up to 10 meters with today's cables. This article helps the designer find an appropriate cable solution based upon actual requirements and the environment.

References

For more information related to this article, you can download an Acrobat Reader file at www-s.ti.com/sc/techlit/*litnumber* and replace "*litnumber*" with the **TI Lit. #** for the materials listed below.

Document Title	TI Lit. #
1. "Performance of LVDS with Different Cables," Application Report	slla053
2. "Low Voltage Differential Signaling (LVDS) Evaluation Module (EVM)," User's Guide	—

Related Web sites

www.ti.com/sc/docs/products/msp/intrface/index.htm
www.ti.com/sc/docs/products/analog/sn65lvds31.html
www.ti.com/sc/docs/products/analog/sn65lvds32.html

Design of op amp sine wave oscillators

By Ron Mancini
Senior Application Specialist, Operational Amplifiers

Criteria for oscillation

The canonical form of a feedback system[1] is shown in Figure 1, and Equation 1 describes the performance of any feedback system (an amplifier with passive feedback components constitutes a feedback system).

Figure 1. Canonical form of a feedback circuit

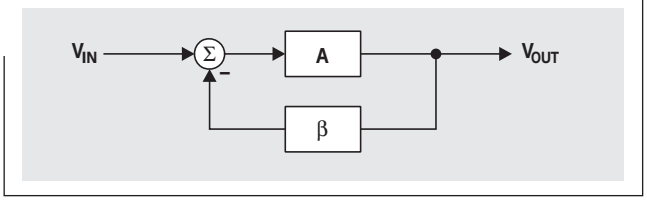

$$\frac{V_{OUT}}{V_{IN}} = \frac{A}{1 + A\beta} \quad (1)$$

Oscillation results from an unstable state; i.e., the feedback system can't find a stable state because its transfer function can't be satisfied. Equation 1 becomes unstable when $(1+A\beta) = 0$ because A/0 is an undefined state. Thus, the key to designing an oscillator is to insure that $A\beta = -1$ (called the Barkhausen criterion), or using complex math the equivalent expression is $A\beta = 1\angle -180°$. The –180° phase shift criterion applies to negative feedback systems, and 0° phase shift applies to positive feedback systems.

The output voltage of a feedback system heads for infinite voltage when $A\beta = -1$. When the output voltage approaches either power rail, the active devices in the amplifiers change gain, causing the value of A to change so the value of $A\beta \neq -1$; thus, the charge to infinite voltage slows down and eventually halts. At this point one of three things can occur. First, nonlinearity in saturation or cutoff can cause the system to become stable and lock up. Second, the initial charge can cause the system to saturate (or cut off) and stay that way for a long time before it becomes linear and heads for the opposite power rail. Third, the system stays linear and reverses direction, heading for the opposite power rail. Alternative two produces highly distorted oscillations (usually quasi square waves), and the resulting oscillators are called relaxation oscillators. Alternative three produces sine wave oscillators.

All oscillator circuits were built with TLV247X op amps, 5% resistors, and 20% capacitors; hence, component tolerances cause differences between ideal and measured values.

Phase shift in oscillators

The 180° phase shift in the equation $A\beta = 1\angle -180°$ is introduced by active and passive components. Like any well-designed feedback circuit, oscillators are made dependent on passive component phase shift because it is accurate and almost drift-free. The phase shift contributed by active components is minimized because it varies with temperature, has a wide initial tolerance, and is device-dependent. Amplifiers are selected such that they contribute little or no phase shift at the oscillation frequency.

A single pole RL or RC circuit contributes up to 90° phase shift per pole, and because 180° is required for oscillation, at least two poles must be used in oscillator design. An LC circuit has two poles; thus, it contributes up to 180° phase shift per pole pair, but LC and LR oscillators are not considered here because low frequency inductors are expensive, heavy, bulky, and non-ideal. LC oscillators are designed in high-frequency applications, beyond the frequency range of voltage feedback op amps, where the inductor size, weight, and cost are less significant. Multiple RC sections are used in low-frequency oscillator design in lieu of inductors.

Phase shift determines the oscillation frequency because the circuit oscillates at the frequency that accumulates –180° phase shift. The rate of change of phase with frequency, $d\phi/dt$, determines frequency stability. When buffered RC sections (an op amp buffer provides high-input and low-output impedance) are cascaded, the phase shift multiplies by the number of sections, n (see Figure 2).

Continued on next page

Figure 2. Phase plot of RC sections

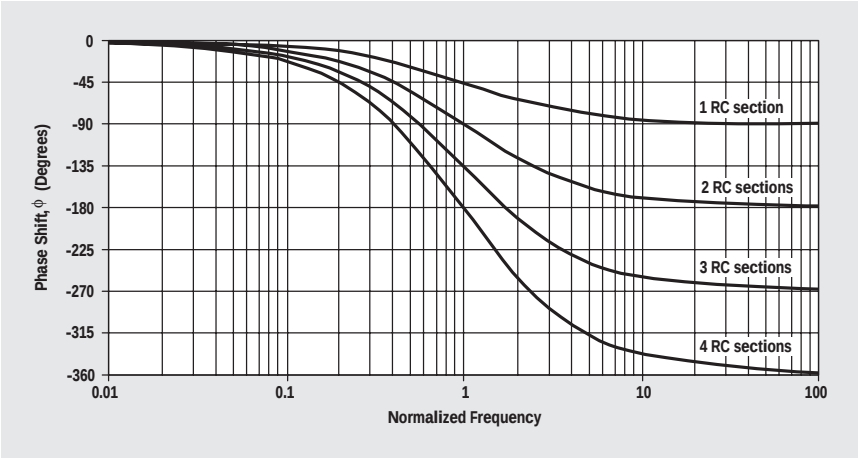

Continued from previous page

Although two cascaded RC sections provide 180° phase shift, dϕ/dt at the oscillator frequency is low, thus oscillators made with two cascaded RC sections have poor frequency stability. Three equal cascaded RC filter sections have a higher dϕ/dt, and the resulting oscillator has improved frequency stability. Adding a fourth RC section produces an oscillator with an excellent dϕ/dt, thus this is the most stable oscillator configuration. Four sections are the maximum number used because op amps come in quad packages, and the four-section oscillator yields four sine waves that are 45° phase shifted relative to each other, so this oscillator can be used to obtain sine/cosine or quadrature sine waves.

Crystal or ceramic resonators make the most stable oscillators because resonators have an extremely high dϕ/dt resulting from their non-linear properties. Resonators are used for high-frequency oscillators, but low-frequency oscillators do not use resonators because of size, weight, and cost restrictions. Op amps are not used with crystal or ceramic resonator oscillators because op amps have low bandwidth. Experience shows that it is more cost-effective to build a high-frequency crystal oscillator and count down the output to obtain a low frequency than it is to use a low-frequency resonator.

Gain in oscillators

The oscillator gain must equal one (Aβ = 1∠–180°) at the oscillation frequency. The circuit becomes stable when the gain exceeds one and oscillations cease. When the gain exceeds one with a phase shift of –180°, the active device non-linearity reduces the gain to one. The non-linearity happens when the amplifier swings close to either power rail because cutoff or saturation reduces the active device (transistor) gain. The paradox is that worst-case design practice requires nominal gains exceeding one for manufacturability, but excess gain causes more distortion of the output sine wave.

When the gain is too low, oscillations cease under worst-case conditions, and when the gain is too high, the output wave form looks more like a square wave than a sine wave. Distortion is a direct result of excess gain overdriving the amplifier; thus, gain must be carefully controlled in low-distortion oscillators. Phase-shift oscillators have distortion, but they achieve low-distortion output voltages because cascaded RC sections act as distortion filters. Also, buffered phase-shift oscillators have low distortion because the gain is controlled and distributed among the buffers.

Some circuit configurations (Wien-bridge) or low-distortion specifications require an auxiliary circuit to adjust the gain. Auxiliary circuits range from inserting a non-linear component in the feedback loop, to automatic gain control (AGC) loops, to limiting by external components.

Figure 3. Wien-bridge circuit schematic

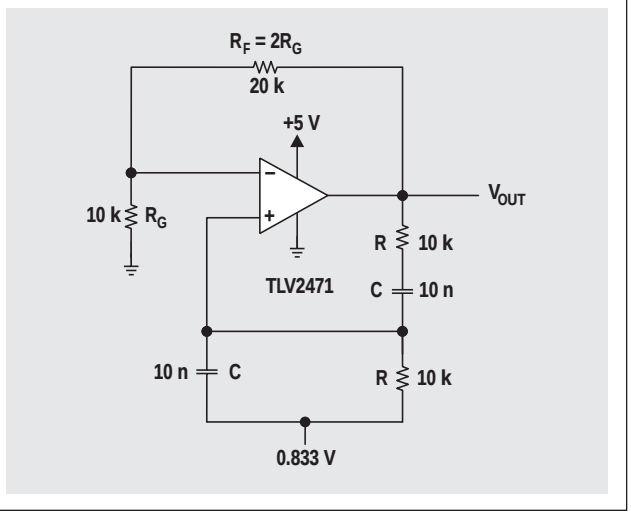

Figure 4. Wien-bridge oscillator with non-linear feedback

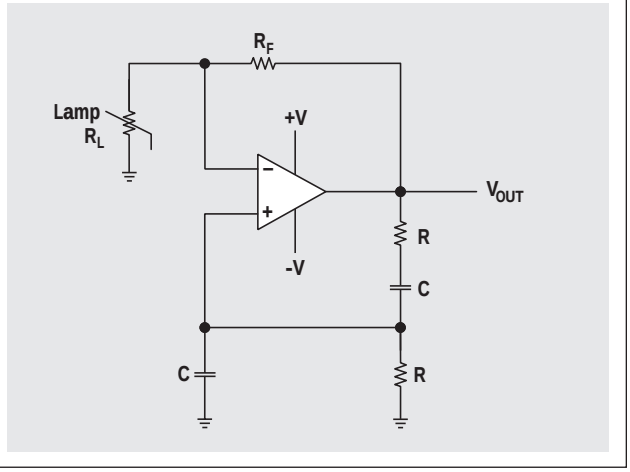

Wien-bridge oscillator

Figure 3 gives the Wien-bridge circuit configuration. The loop is broken at the positive input, and the return signal is calculated in Equation 2 below.

$$\frac{V_{RETURN}}{V_{OUT}} = \frac{\frac{R}{RCs+1}}{\frac{R}{RCs+1}+R+\frac{1}{Cs}} = \frac{1}{3+RCs+\frac{1}{RCs}} = \frac{1}{3+j\left(RC\omega-\frac{1}{RC\omega}\right)}, \quad (2)$$

where s = jω and j = $\sqrt{-1}$.

When ω = 2πf = 1/RC, the feedback is in phase (this is positive feedback), and the gain is 1/3, so oscillation requires an amplifier with a gain of 3. When $R_F = 2R_G$, the amplifier gain is 3 and oscillation occurs at f = 1/2πRC. The circuit oscillated at 1.65 kHz rather than 1.59 kHz with the component values shown in Figure 3, but the distortion

is noticeable. Figure 4 shows a Wien-bridge circuit with non-linear feedback. The lamp resistance, R_L, is nominally selected as half the feedback resistance, R_F, at the lamp current established by R_F and R_L. The non-linear relationship between the lamp current and resistance keeps output voltage changes small.

Some circuits use diode limiting in place of a non-linear feedback component. The diodes reduce the distortion by providing a soft limit for the output voltage. AGC must be used when neither of these techniques yields low distortion. A typical Wien-bridge oscillator with an AGC circuit is shown in Figure 5.

The negative sine wave is sampled by D_1, and the sample is stored on C_1. R_1 and R_2 are chosen to center the bias on Q_1 so that $(R_G + R_{Q1}) = R_F/2$ at the desired output voltage. When the output voltage drifts high, Q_1 increases resistance, thus decreasing the gain. In the oscillator shown in Figure 3, the 0.833-volt power supply is applied to the positive op amp input to center the output quiescent voltage at $V_{CC}/2 = 2.5$ V.

Figure 5. Wien-bridge oscillator with AGC

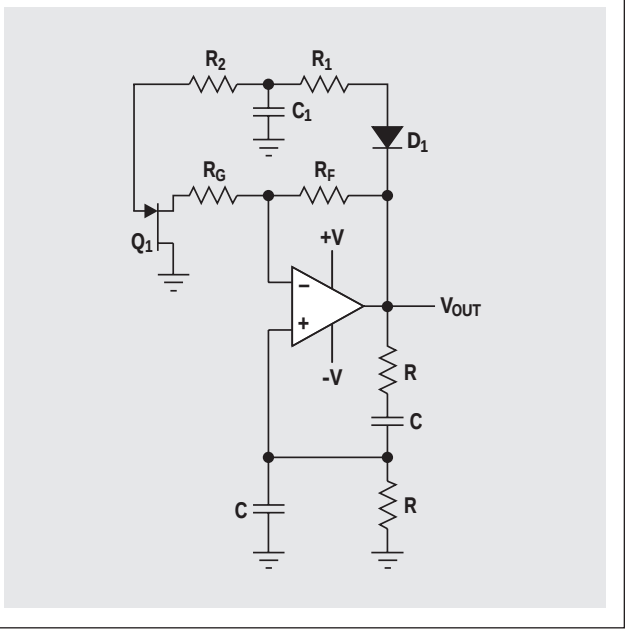

Phase-shift oscillator (one op amp)

A phase-shift oscillator can be built with one op amp as shown in Figure 6.

The normal assumption is that the phase-shift sections are independent of each other. Then Equation 3 is written:

$$A\beta = A\left(\frac{1}{RCs+1}\right)^3 \qquad (3)$$

The loop phase shift is –180° when the phase shift of each section is –60°, and this occurs when $\omega = 2\pi f = 1.732/RC$ because the tangent 60° = 1.73. The magnitude of β at this point is $(1/2)^3$, so the gain, A, must be equal to 8 for the system gain to be equal to 1.

The oscillation frequency with the component values shown in Figure 6 is 3.76 kHz rather than the calculated oscillation frequency of 2.76 kHz. Also, the gain required to start oscillation is 26 rather than the calculated gain of 8. These discrepancies are partially due to component variations, but the biggest contributing factor is the incorrect assumption that the RC sections do not load each other. This circuit configuration was very popular when active components were large and expensive, but now op amps are inexpensive and small and come four in a package, so the single op amp phase-shift oscillator is losing popularity.

Continued on next page

Figure 6. Phase-shift oscillator (one op amp)

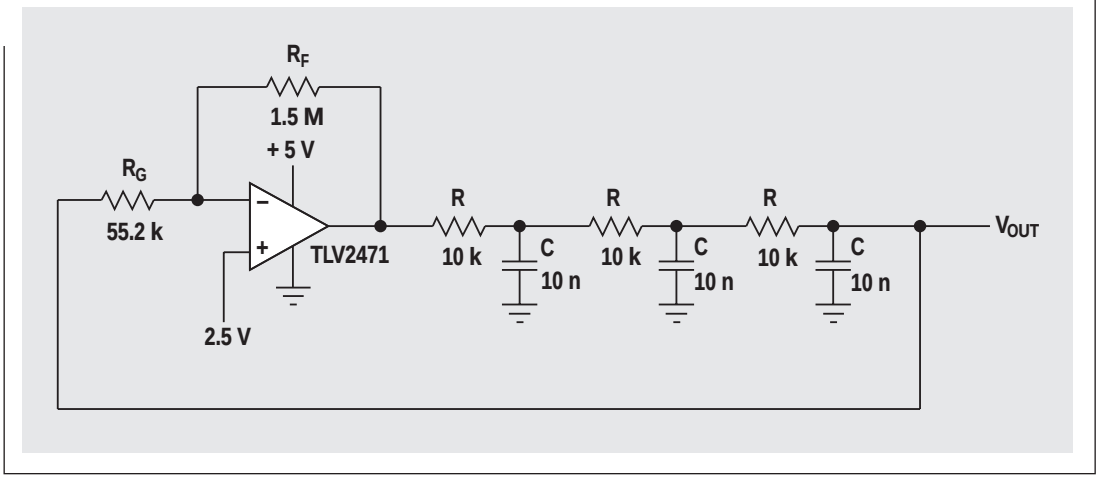

Figure 7. Buffered phase-shift oscillator

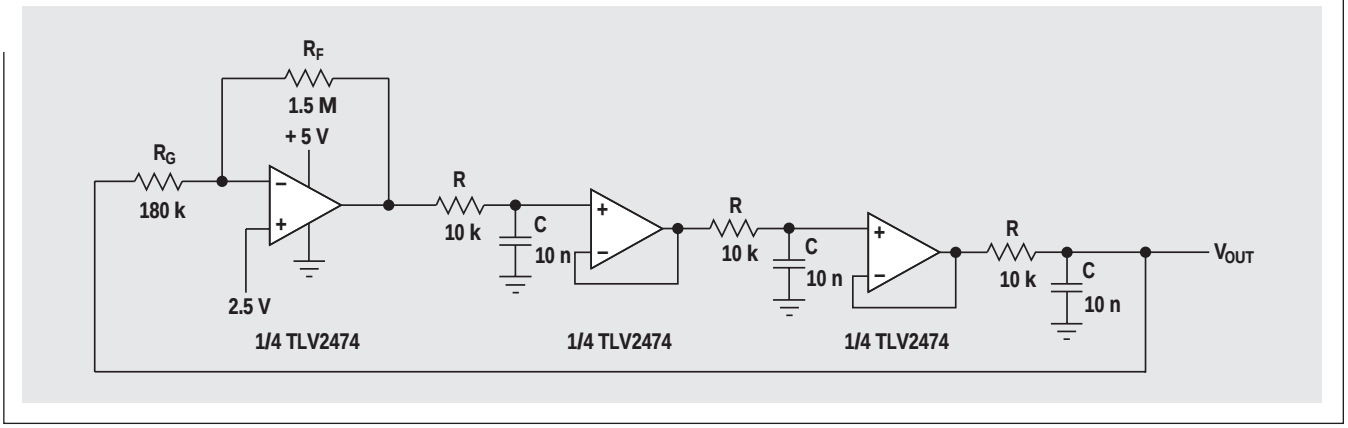

Continued from previous page

Buffered phase-shift oscillator

The buffered phase-shift oscillator shown in Figure 7 oscillated at 2.9 kHz compared to an ideal frequency of 2.76 kHz, and it oscillated with a gain of 8.33 compared to an ideal gain of 8.

The buffers prevent the RC sections from loading each other, hence the buffered phase-shift oscillator performs closer to the calculated frequency and gain. The gain setting resistor, R_G, loads the third RC section, and if the fourth op amp in a quad op amp buffers this RC section, the performance becomes ideal. Low-distortion sine waves can be obtained from either phase-shift oscillator, but the purest sine wave is taken from the output of the last RC section. This is a high-impedance node, so a high-impedance input is mandated to prevent loading and frequency shifting with load variations.

Quadrature oscillator

The quadrature oscillator is another type of phase-shift oscillator, but the three RC sections are configured so that each section contributes 90° of phase shift. The outputs are labeled sine and cosine (quadrature) because there is a 90° phase shift between op amp outputs (see Figure 8). The loop gain is calculated in Equation 4.

$$A\beta = \left(\frac{1}{R_1C_1s}\right)\left(\frac{R_3C_3s+1}{R_3C_3s(R_2C_2s+1)}\right) \quad \textbf{(4)}$$

When $R_1C_1 = R_2C_2 = R_3C_3$, Equation 4 reduces to Equation 5.

$$A\beta = \frac{1}{(RCs)^2} \quad \textbf{(5)}$$

When $\omega = 1/RC$, Equation 5 reduces to $1\angle -180°$, so oscillation occurs at $\omega = 2\pi f = 1/RC$. The test circuit oscillated at 1.65 kHz rather than the calculated 1.59 kHz, and the discrepancy is attributed to component variations.

Figure 8. Quadrature oscillator

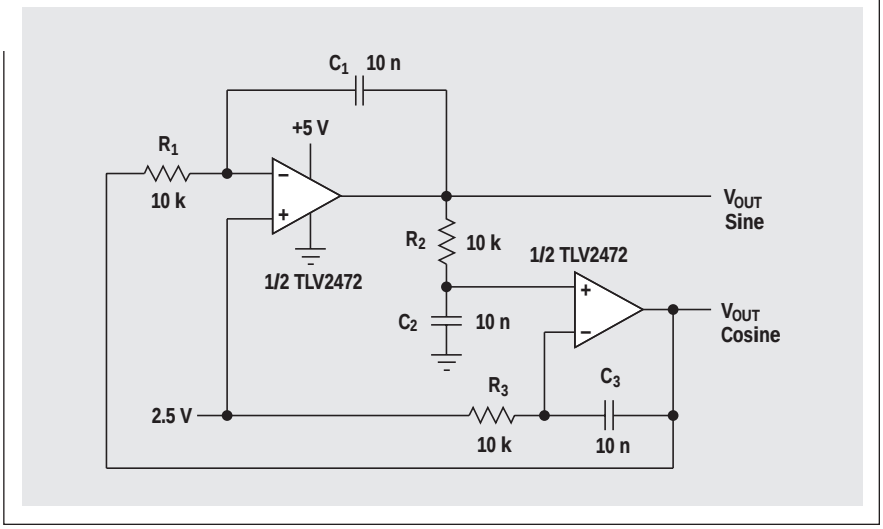

Bubba oscillator

The Bubba oscillator (Figure 9) is another phase-shift oscillator, but it takes advantage of the quad op amp package to yield some unique advantages. Four RC sections require 45° phase shift per section, so this oscillator has an excellent dϕ/dt to minimize frequency drift. The RC sections each contribute 45° phase shift, so taking outputs from alternate sections yields low-impedance quadrature outputs. When an output is taken from each op amp, the circuit delivers four 45° phase-shifted sine waves. The loop equation is:

$$A\beta = A\left(\frac{1}{RCs+1}\right)^4 \qquad (6)$$

When ω = 1/RCs, Equation 6 reduces to Equations 7 and 8.

$$|\beta| = \left|\left(\frac{1}{1+j}\right)^4\right| = \frac{1}{\sqrt{2}^4} = \frac{1}{4} \qquad (7)$$

$$\text{Phase} = \text{Tan}^{-1}1 = 45° \qquad (8)$$

The gain, A, must equal 4 for oscillation to occur. The test circuit oscillated at 1.76 kHz rather than the ideal frequency 1.72 kHz when the gain was 4.17 rather than the ideal gain of 4. With low gain, A, and low bias current op amps, the gain setting resistor, R_G, does not load the last RC section thus insuring oscillator frequency accuracy. Very low-distortion sine waves can be obtained from the junction of R and R_G. When low-distortion sine waves are required at all outputs, the gain should be distributed between all the op amps. The non-inverting input of the gain op amp is biased at 0.5 V to set the quiescent output voltage at 2.5 V. Gain distribution requires biasing of the other op amps, but it has no effect on the oscillator frequency.

Figure 9. Bubba oscillator

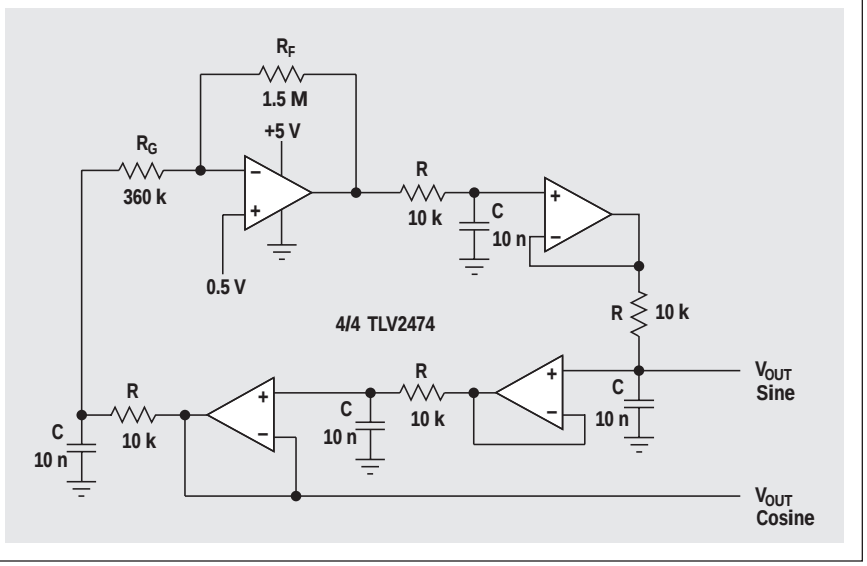

Summary

Op amp oscillators are restricted to the lower end of the frequency spectrum because op amps do not have the required bandwidth to achieve low phase shift at high frequencies. The new current feedback op amps are very hard to use in oscillator circuits because they are sensitive to feedback capacitance. Voltage feedback op amps are limited to a few hundred kHz because they accumulate too much phase shift.

The Wien-bridge oscillator has few parts, and its frequency stability is good. Taming the distortion in a Wien-bridge oscillator is harder than getting the circuit to oscillate. The quadrature oscillator only requires two op amps, but it has high distortion. Phase-shift oscillators, especially the Bubba oscillator, have less distortion coupled with good frequency stability. The improved performance of the phase-shift oscillators comes at a cost of higher component count.

References

For more information related to this article, you can download an Acrobat Reader file at www-s.ti.com/sc/techlit/*litnumber* and replace *"litnumber"* with the **TI Lit. #** for the materials listed below.

Document Title	**TI Lit. #**
1. "Feedback Amplifier Analysis Tools"	.sloa017

Related Web sites

www.ti.com/sc/amplifiers
www.ti.com/sc/docs/products/analog/tlv2471.html
www.ti.com/sc/docs/products/analog/tlv2472.html
www.ti.com/sc/docs/products/analog/tlv2474.html

Fully differential amplifiers

By Jim Karki
Systems Specialist, High-Speed Amplifiers

Introduction

Professional audio engineers usually use the term "balanced" to refer to differential signal transmission. This imparts the idea of symmetry, which is very important in differential systems. The driver has balanced outputs, the line has balanced characteristics, and the receiver has balanced inputs.

There are two methods commonly used to manipulate differential signals: electronic and transformer.

1. Electronic methods have advantages such as low cost, small size and weight, and superior frequency response at low frequency and dc.
2. The advantages that transformers offer are excellent common-mode rejection ratio (CMRR), galvanic isolation, no power consumption (efficiencies near 100%), and immunity to very hostile EMC environments.

This article focuses on integrated, fully differential amplifiers for signal conditioning differential signals. Basic operations, such as how to transform single-ended signals into differential signals and how to construct active anti-alias filters, are discussed.

What is an integrated, fully differential amplifier?

An integrated, fully differential amplifier is very similar in architecture to a standard operational amplifier.

Figure 1 shows a simplified version of an integrated, fully differential amplifier. Q1 and Q2 are the input differential pair. In a standard op amp, output current is taken from only one side of the input differential pair and is used to develop a single-ended output voltage. In a fully differential amplifier, currents from both sides are used to develop voltages at the high-impedance nodes formed at the collectors of Q3/Q5 and Q4/Q6. These voltages are then buffered to the differential outputs OUT+ and OUT–.

For a first-order approximation, voltage common to IN+ and IN– does not produce a change in the current flow through Q1 or Q2 and thus produces no output voltage; it is rejected. The output common-mode voltage is not controlled by the input. The V_{CM} error amplifier controls the output common-mode voltage by sampling it, comparing it to the voltage at V_{CM}, and adjusting the internal feedback.

The two complementary amplifier paths share the same input differential pair, their characteristics are very well

Figure 1. Simplified, fully differential amplifier

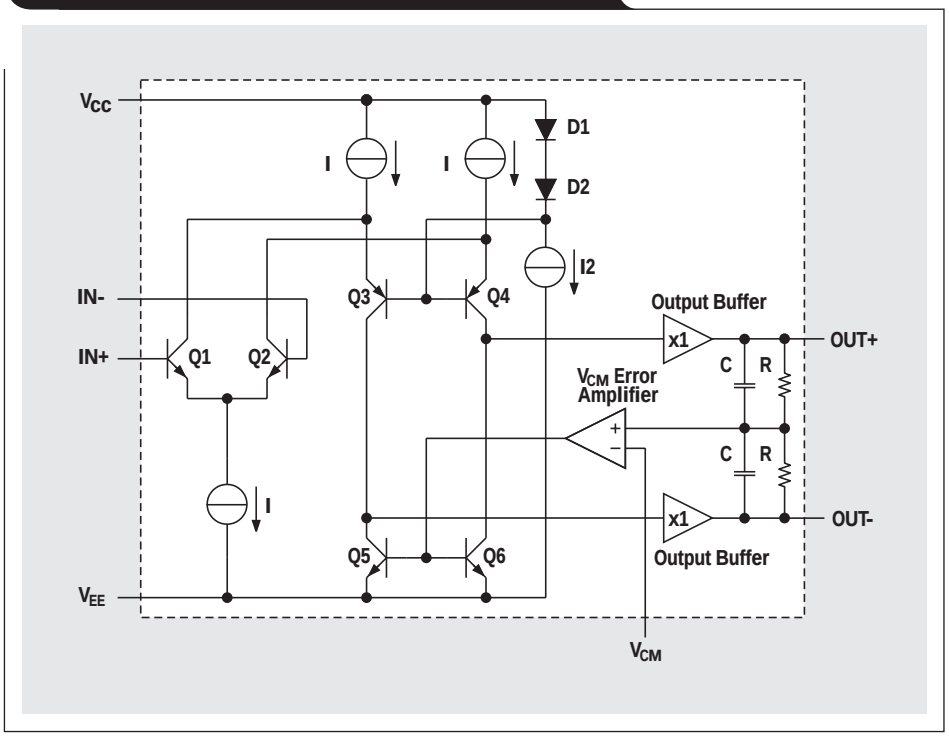

matched, and the architecture keeps their operating points very close to each other. Therefore, distortion in the amplifiers is also matched, resulting in symmetrical distortion of the differential signal. Symmetrical distortions tend to cancel even-order harmonics. Lab testing shows that the second harmonic levels in a differential output are reduced by approximately 5 dB in the THS4141 at 1 MHz when measured differentially as compared to measuring either single-ended output. The measured level of the third harmonic is unchanged.

Voltage definitions

To understand how a fully differential amplifier behaves, it is important to understand the voltage definitions that are used to describe the amplifier. Figure 2 shows a block diagram that represents a fully differential amplifier and its input and output voltage definitions.

The voltage difference between the plus and minus inputs is the input differential voltage, V_{ID}. The average of the two input voltages is the input common-mode voltage, V_{IC}.

The difference between the voltages at the plus and minus outputs is the output differential voltage, V_{OD}. The output common-mode voltage, V_{OC}, is the average of the two output voltages and is controlled by the voltage at V_{CM}.

A_f is the frequency-dependent differential gain of the amplifier, so that $V_{OD} = V_{ID} \times A_f$.

Increased noise immunity

Invariably, when signals are routed from one place to another, noise is coupled into the wiring. In a differential system, keeping the transport wires as close as possible to one another makes the noise coupled into the conductors appear as a common-mode voltage. Noise that is common to the power supplies will also appear as a common-mode voltage. Since the differential amplifier rejects common-mode voltages, the system is more immune to external noise. Figure 3 shows the noise immunity of a fully differential amplifier.

Increased dynamic range

Due to the change in phase between the differential outputs, the dynamic range increases by 2x over a single-ended output with the same voltage swing (see Figure 4).

Continued on next page

Figure 2. Fully differential amplifier voltage definitions

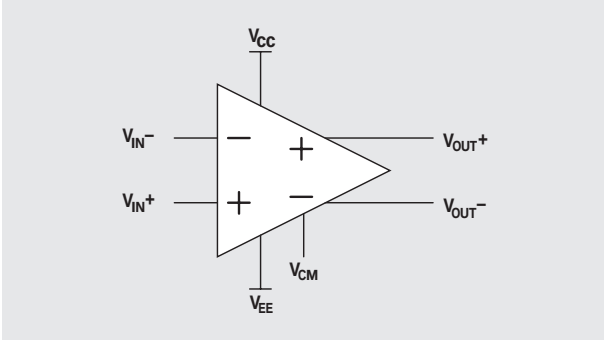

Input voltage definition

$$V_{ID} = (V_{IN}+) - (V_{IN}-) \qquad V_{IC} = \frac{(V_{IN}+) + (V_{IN}-)}{2}$$

Output voltage definition

$$V_{OD} = (V_{OUT}+) - (V_{OUT}-) \qquad V_{OC} = \frac{(V_{OUT}+) + (V_{OUT}-)}{2}$$

Transfer function $V_{OD} = V_{ID} \times A_f$

Output common-mode voltage $V_{OC} = V_{CM}$

Figure 3. Fully differential amplifier noise immunity

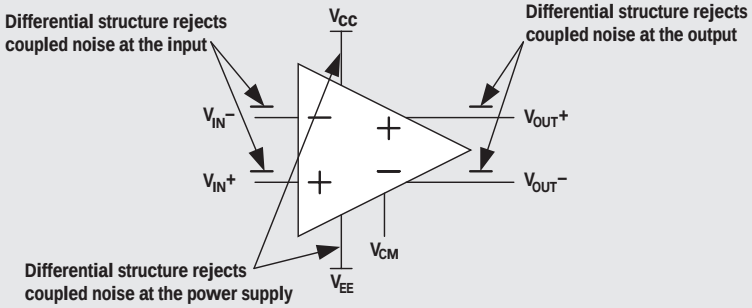

Figure 4. Differential output voltage swing

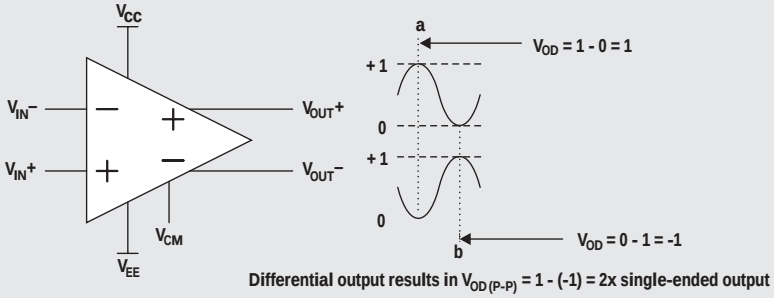

Continued from previous page

Basic circuits

To maintain balance in a fully differential amplifier, symmetrical feedback must be taken from both outputs and applied to both inputs. The two sides form symmetrical inverting amplifiers, and inverting op amp topologies are easily adapted to fully differential amplifiers. Figure 5 shows how to maintain a balanced amplifier by using symmetrical feedback, where the feedback resistors, R_F, and the input resistors, R_G, are equal.

It is important to maintain symmetry in the two feedback paths to maintain good CMRR performance. CMRR is directly proportional to the resistor matching error. For example, a 0.1% error results in 60 dB of CMRR. For small variations in feedback due to mismatched resistors, the differential gain of the amplifier is approximately the average gain of the two sides. Output balance is maintained by the V_{CM} error amplifier.

In the past, generation of differential signals has been cumbersome. Different means have been used, requiring as many as three amplifiers and dc blocking capacitors to set the output common-mode voltage. The integrated, fully differential amplifier provides a more elegant solution. Figure 6 shows an example of converting single-ended signals to differential signals.

Active anti-alias filtering

A major application for fully differential amplifiers is signal conditioning ADC inputs. Low-pass filters are needed to keep high-frequency noise from aliasing into the frequency band of interest. Multiple feedback (MFB) is a good topology that is adapted easily to a fully differential amplifier. An MFB circuit is used to realize one complex pole pair in the transfer function of a second-order low-pass filter. An example is shown in Figure 7.

Figure 5. Amplifying differential signals

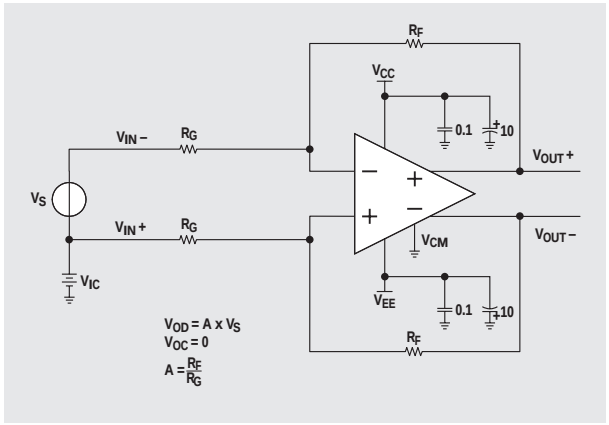

Figure 6. Converting single-ended signals to differential signals

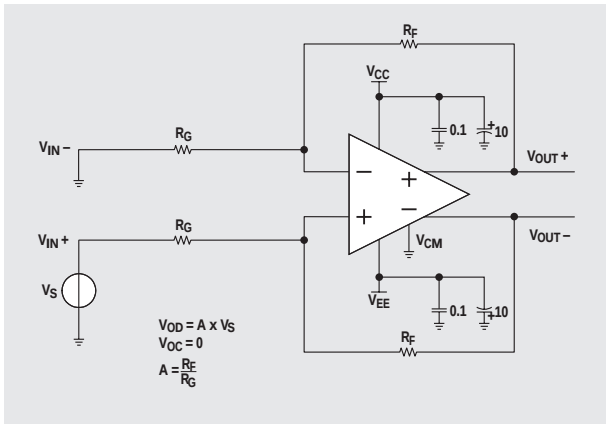

Figure 7. Low-pass, fully differential filter driving an ADC

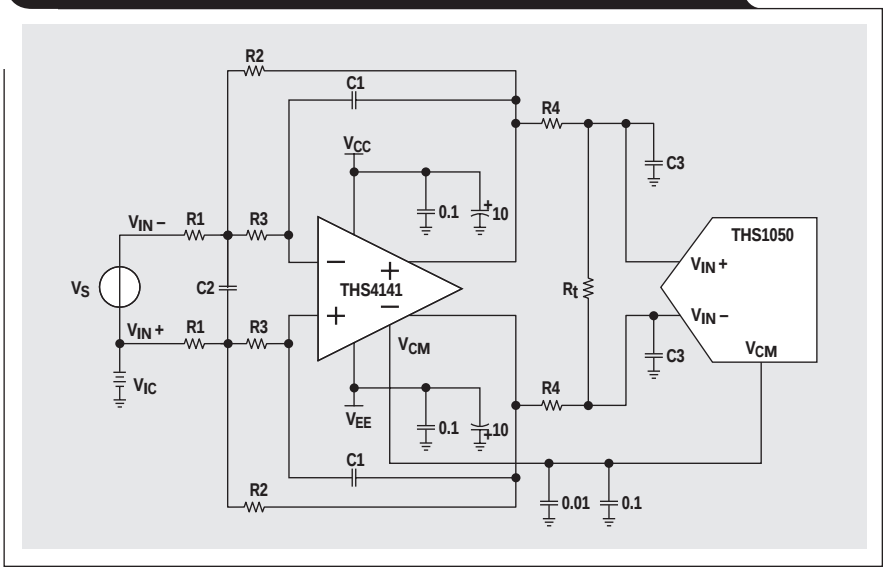

The transfer function for this filter circuit is:

$$H_d(f) = \left[\frac{K}{-\left(\frac{f}{FSF \times f_C}\right)^2 + \frac{1}{Q}\frac{jf}{FSF \times f_C} + 1} \right] \times \left(\frac{\frac{R_t}{2R4 + R_t}}{1 + \frac{j2\pi f R4 R_t C3}{2R4 + R_t}} \right),$$

where $K = \frac{R2}{R1}$, $FSF \times f_C = \frac{1}{2\pi\sqrt{2 \times R2R3C1C2}}$, and

$$Q = \frac{\sqrt{2 \times R2R3C1C2}}{R3C1 + R2C1 + KR3C1}.$$

K sets the pass-band gain, f_C is the cut-off frequency for the filter, FSF is a frequency scaling factor, and Q is the quality factor.

$$FSF = \sqrt{Re^2 + |Im|^2}, \quad \text{and } Q = \frac{\sqrt{Re^2 + |Im|^2}}{2Re},$$

where Re is the real part, and Im is the imaginary part of the complex pole pair. Setting R2=R, R3=mR, C1=C, and C2=nC results in:

$$FSF \times f_C = \frac{1}{2\pi RC\sqrt{2 \times mn}}, \quad \text{and } Q = \frac{\sqrt{2 \times mn}}{1 + m(1 - K)}.$$

Start by determining the ratios, m and n, required for the gain and Q of the filter type being designed, then select C, and calculate R for the desired f_C.

The combination of R4, R_t, and C3 has multiple effects. R4 isolates the amplifier output from the input of the ADC. R4 and Rt provide for double termination of the transmission line between the amplifier and the ADC, and form a voltage divider. C3 helps absorb charge injection from the ADC's input. R4 and C3 form a real pole that can be used to make a third-order filter, in conjunction with the complex pole pair from the MFB stage, or it can simply be placed above the frequencies of interest.

The proper V_{CM} is provided as an output by some ADCs with differential inputs. Typically, all that needs to be done is to provide bypass capacitors—0.1 µF and/or 0.01 µF. If not provided, V_{CM} can be generated from the ADC's reference voltages as shown in Figure 8. The voltage at the summing node will be the midpoint between the reference voltage and will center V_{OC} in the middle of the ADC's input range.

Figure 8. Typical V_{CM} error source

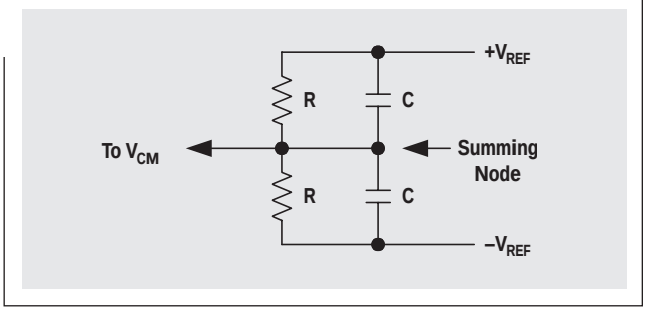

Each power pin should have a 6.8-µF to 10-µF tantalum capacitor in parallel with a 0.01-µF to 0.1-µF ceramic capacitor located very close by. Figure 7 shows 10-µF and 0.1-µF power-supply bypass capacitors.

Conclusion

Integrated, fully differential amplifiers are very similar to standard single-ended op amps except that output is taken from both sides of the input differential pair to produce a differential output.

Differential systems provide increased immunity to external noise, reduced even-order harmonics, and twice the dynamic range when compared to single-ended systems.

Inverting amplifier topologies are adapted easily to fully differential amplifiers by implementing two symmetric feedback paths.

Integrated, fully differential amplifiers are well-suited for driving differential ADC inputs. They provide an easy means for anti-alias filtering, and the required common-mode voltage is set easily via the V_{CM} input.

Related Web sites

www.ti.com/sc/amplifiers
www.ti.com/sc/docs/products/analog/ths4141.html
www.ti.com/sc/docs/products/analog/ths1050.html

The PCB is a component of op amp design

By Bruce Carter
Senior Applications Specialist

Most analog designers are familiar with how to use ICs and passive components to implement a design. There is one additional circuit component, however, that must be considered for the design to be a success—the printed circuit board on which the circuit is to be located. Analog circuitry is very different from digital circuitry and should be partitioned in its own section of the board with special layout techniques. Printed circuit board effects become most apparent in high-speed analog circuits, but common mistakes described later can affect even the performance of audio circuits. Any effect caused by the PCB itself should be minimized, such that the operation of the analog circuitry in production will be the same as the performance of the design and prototype.

Normal design cycles, particularly of large digital boards, dictate layout of the PCB as soon as possible. The digital circuitry may have been simulated, but in most cases the production PCB itself is the prototype and may even be sold to a customer. By implementing "cuts and jumpers" and reprogramming gate arrays or flash memories, digital designers can correct small mistakes and go on with the next project. This is not the case with analog circuitry. Some common analog designs cannot be corrected by the "cut and jumper" method. They can and will render the entire PCB unusable. It is very important for the digital designer, who is used to "cuts and jumpers," to read and understand this article prior to releasing a board to a layout service. A little extra thought during development can save a board worth thousands of dollars from becoming "scrap" because of blunders in a tiny section of analog circuitry.

Noise sources

Noise is conducted into the analog circuit through its connections to other circuits. Analog circuitry must be connected to the "outside world" by at least a ground connection, a power connection, an input, and an output. Noise can be conducted into the circuit through all of these paths, as well as through any others that are present.

Noise is radiated into the analog circuitry from many external sources. The most common is high-speed digital logic, including DSP chips that reside in the system or on the same PCB. High-speed clocks and switching digital signals create considerable radio frequency interference (RFI). Other sources of radiated noise include a switching power supply, cellular telephones, broadcast radio and TV, fluorescent lighting, nearby PCs, and lightning in thunderstorms. Even if the analog circuitry only operates at audio frequencies, RFI may produce noticeable noise in the output.

PCB materials—choose the right one for your application

PC board materials are available in various grades, as defined by the National Electrical Manufacturers Association (NEMA). It would be very convenient for designers if this organization were closely allied with the electronics industry, controlling parameters such as resistivity and dielectric constant of the material. Unfortunately, that is not the case. NEMA is an electrical safety organization, and the different PCB grades primarily describe the flammability, high-temperature stability, and moisture absorption of the board. Therefore, specifying a given NEMA grade does not guarantee electrical parameters of the material. If this becomes critical, consult the manufacturer of the raw board stock.

Raw PCB stock is graded in flammability ratings (FR) from 1 to 5, with 1 being the most flammable and 5 being the least. FR-4 is commonly used in industrial quality equipment, while FR-2 is used in high-volume consumer applications. Although there are no set rules for this, it appears to be an industry "standard." Deviating from it without good reason can limit the number of suppliers of raw board material and the number of PCB houses that can fabricate the board, since their tooling is already set up for these materials.

In selecting a board material, pay careful attention to the moisture absorption. Just about every desirable performance characteristic of the board will be negatively impacted by moisture. This includes surface resistance of the board, dielectric leakage, high-voltage breakdown and arcing, and mechanical stability. Also pay attention to the operating temperature. High operating temperatures can occur in unexpected places, such as in proximity to large digital chips that are switching at high speeds. Heat rises, so be aware that if one of those 300-pin ICs is located directly under a sensitive analog circuit, both the PCB and the circuit characteristics may vary with the temperature.

How many layers are best?

Many times, the number of board layers has already been determined by system constraints. If the designer has a choice, however, there are some guidelines.

Very simple consumer electronics are sometimes fabricated on single-sided PCBs, keeping the raw board material inexpensive (FR-1 or FR-2) with thin copper cladding. These designs frequently include many jumper wires, simulating the circuit routing on a double-sided board.

This technique is only recommended for low-frequency circuitry. For reasons described later, this type of design is extremely susceptible to radiated noise. Therefore, it is actually more complex to design a board of this type due to the many things that can go wrong.

The next level of complexity is double-sided. Initially, this type of board would seem to lend itself to easier routing because it has two layers of foil, and it is possible to route signals by crossing traces on different layers. While that is certainly possible, it is not recommended for analog circuitry. Wherever possible, the bottom layer should be devoted to a continuous ground plane, and all other signals routed on the top layer. A ground plane provides several benefits:

- Ground is frequently the most common connection in the circuit. Having it continuous on the bottom layer usually makes the most sense for circuit routing.
- It increases the mechanical strength of the board.
- It lowers the impedance of all ground connections in the circuit, which reduces undesirable conducted noise.
- It adds a distributed capacitance to every net in the circuit, helping to suppress radiated noise.
- It acts as a shield to radiated noise coming from underneath the board.

Double-sided boards, in spite of their benefits, are not the best method of construction, especially for sensitive or high-speed designs. The most common board thickness is 1.5 mm. This separation is too great for full realization of some of the benefits just listed. Distributed capacitance, for example, is very low due to the separation. Critical designs call for multi-layer boards. Although more expensive, they provide the following benefits:

- They have better routing for power as well as ground connections. If the power is also on a plane, it is available to all points in the circuit simply by adding vias.
- Other layers are available for signal routing, making routing easier.
- There will be distributed capacitance between the power and ground planes, reducing high-frequency noise.

The decision to use multi-layer boards is complex, requiring the designer to weigh board cost against performance. Be sure to consider the cost of qualification testing, if any. Multi-layer boards present a much lower design risk.

Grounding

Good grounding is a system-level design consideration. Proper grounding should be planned into the product from the first conceptual design reviews.

Separate grounding for analog and digital portions of circuitry is one of the simplest and most effective methods for noise suppression. One or more layers on multi-layer PCBs usually are devoted to ground planes. If the designer is not careful, the analog circuitry will be connected directly to these “ground” planes. The analog circuitry return, after all, is the same net in the netlist as digital return. Auto-routers respond accordingly, connecting all of the grounds together—creating a disaster.

Ground and power planes are at the same AC potential, due to decoupling capacitors and distributed capacitance. Therefore, it is important to isolate the power planes as well.

Do not overlap digital and analog planes (see Figure 1). Place analog power coincident with analog ground, and digital power coincident with digital ground. If any portion of analog and digital planes overlaps, the distributed capacitance between the overlapping portions will couple high-speed digital noise into the analog circuitry. This defeats the purpose of isolated planes.

“Separate grounds” does not mean that the grounds are electrically separate in the system. They have to be common at some point—preferably a single, low-impedance point. In the system there is only one ground—the electrical safety ground in an ac-powered system or the battery ground in a dc-powered system. Everything else “returns” to that ground. Refer to everything that is not a ground as a “return.” All returns should be connected together at a single point, which is system “ground.” At some point, this will be the chassis. It is important to avoid ground loops by multiple connections to the chassis (Figure 2). Insuring only one chassis ground point is one of the most difficult aspects of system design.

Continued on next page

Figure 1. Ground and power plane overlap

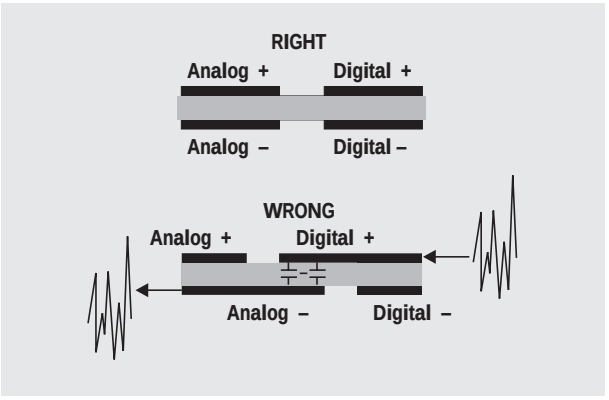

Figure 2. Single ground point minimizes ground loops

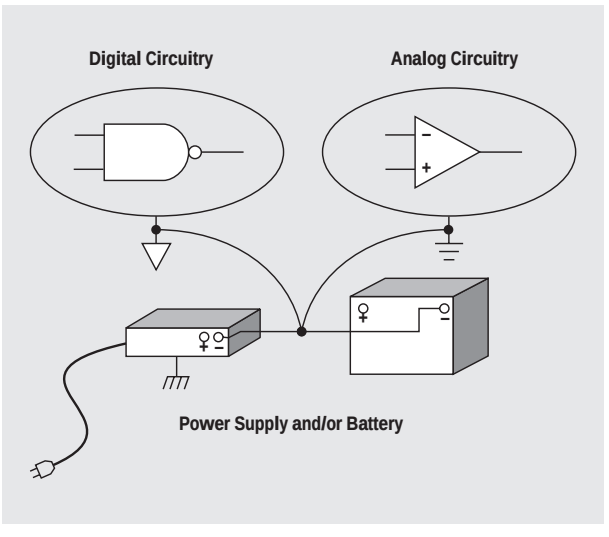

Continued from previous page

If at all possible, dedicate separate connector pins to separate returns, and combine the returns only at system ground. Aging and repeated mating cause connector pins to increase in contact resistance, so several pins are needed. Many digital boards consist of many layers and hundreds or thousands of nets. The addition of one more net is seldom an issue, but the addition of several connector pins almost always is. If this cannot be done, then it will be necessary to make the two returns a single net on the PCB—with very special routing precautions.

It is important to keep digital signals away from analog portions of the circuit. It makes little sense to isolate planes, keep analog traces short, and place passive components carefully if there are high-speed digital traces running right next to the sensitive analog traces. Digital signals must be routed around analog circuitry and not overlap analog ground and power planes (Figure 3).

Most digital clocks are high enough in frequency that even small capacitances between traces and planes can couple significant noise. Remember that it is not only the fundamental frequency of the clock that can cause a potential problem but also the higher-frequency harmonics.

It is a good idea to locate analog circuitry as close as possible to the I/O connections of the board. Digital designers used to high-current ICs will be tempted to make a 50-mil trace run several inches to the analog circuitry—thinking that reducing the resistance in the trace will help get rid of noise. What will actually result is a long, skinny capacitor that couples noise from digital ground and power planes into the op amp, making the problem worse!

The frequency characteristics of passive components

Most designers are totally ignorant of the frequency limitations of the passive components they use in analog circuitry. Passive components have limited frequency ranges, and operation of the part outside of that range can have some very unexpected results. In most cases, a “right” passive component will fit on the same pads as a “wrong” passive component, but not always. Start the design process by carefully considering the high-frequency characteristics of passive components and putting the correct part outline on the board from the start.

Figure 3. Digital and analog circuitry should not overlap

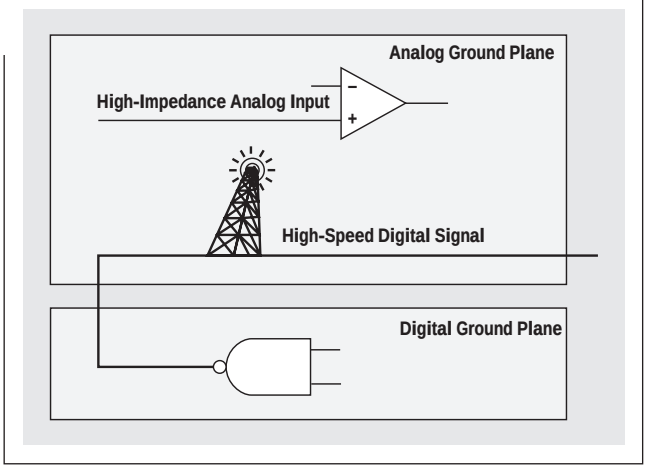

Resistors

High-frequency performance of resistors is approximated by Figure 4.

Figure 4. Resistor characteristics at high frequencies

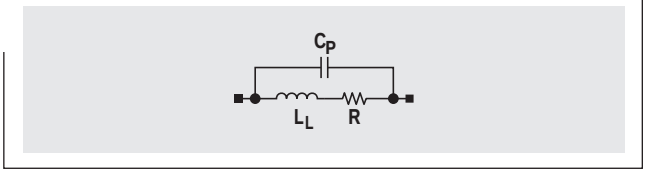

Resistors are constructed three ways—wire wound, carbon composition, and film. It does not take a lot of imagination to understand how wire wound resistors can become inductive, because they are coils of resistive wire. Film resistors are coils of thin metallic film. Therefore, they are also inductive at high frequencies.

The end caps of resistors are parallel, and there will be an associated capacitance. Usually, the resistance will make the parasitic capacitor so “leaky” that the capacitance does not matter. For very high resistances, the capacitance will appear in parallel with the resistance, lowering its impedance at high frequencies.

Capacitors

High-frequency performance of capacitors is approximated by Figure 5.

Figure 5. Capacitor characteristics at high frequencies

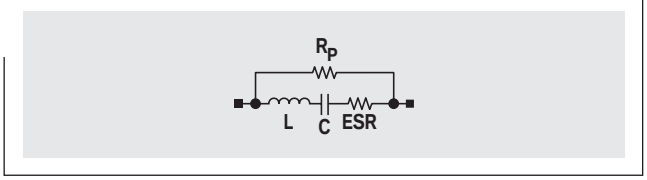

Film and electrolytic capacitors have layers of material wound around each other, which creates a parasitic inductance. Self-inductance effects of ceramic capacitors are much smaller, giving them a higher operating frequency. There is also some leakage current from plate to plate R_P, which appears as a resistance in parallel with the capacitor. The most important parasitic component in a capacitor is the equivalent series resistance (ESR). It is due to resistance within the plates and electrolyte of an electrolytic capacitor. Capacitors used for decoupling should be low ESR types, as any series resistance limits the effectiveness of the capacitor for ripple and noise rejection. Elevated temperatures severely increase ESR and can be permanently destructive to capacitors.

The leads of through-hole parts also add a parasitic inductance. For small values of capacitance, it is important to keep the lead lengths short. The combination of parasitic inductance and capacitance can produce resonant circuits.

PCB trace antennas

A board is vulnerable to radiated interference because the pattern of traces and component leads forms antennas. Common wires and PCB traces have inductance that varies between 6 and 12 nH per centimeter. At frequencies above 100 kHz, most PCB traces are inductive, not resistive.

A rule of thumb for antennas is that they begin to couple significant energy at about 1/20 of the wavelength of the received signal. Therefore, a 10-cm trace will begin to be a fairly good antenna at frequencies above 150 MHz. Remember that although the clock generator on a digital PCB may not be operating at a frequency as high as 150 MHz, it approximates a square wave. Square waves will have harmonics throughout the frequency range where PCB conductors become efficient antennas.

A loop also can form an antenna. Most digital designers familiar with aspects of loop antenna theory know not to make loops in critical signal pathways. Some designers who would never think of making a loop with a high-speed clock or reset signal, however, will create a loop by the technique they use for layout of the analog section of the board. Loop antennas constructed as loops of wire are easy to visualize. Slot antennas are harder to visualize, but just as efficient. Consider the three cases illustrated in Figure 6.

Version A is a bad design. It does not utilize an analog ground plane at all. A loop is formed by the ground and signal traces. An electric field, E, and perpendicular magnetic field, H, are created and form the basis of a loop antenna.

Version B is a better design, but there is intrusion into the ground plane, presumably to make room for a signal trace. A much smaller slot antenna is formed by the difference in pathways between signal and return.

Version C is the best design. Signal and return are coincident with each other, eliminating loop antenna effects completely. Note that there are cutouts for the ICs, but they are located away from the return path for the signal.

Figure 6. Design comparisons

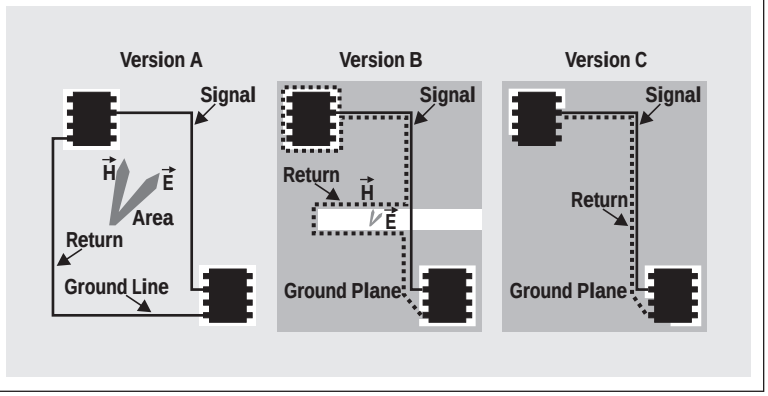

When a PCB trace turns a corner at a 90° angle, a reflection can occur. This is primarily due to the change of width of the trace. At the apex of the turn, the trace width is increased to 1.414 times its normal width (see Figure 7). This upsets the transmission line characteristics, especially the distributed capacitance and self-inductance of the trace, resulting in the reflection. It is a given that not all PCB traces can be straight, so they will have to turn corners. Most CAD systems give some rounding effect on the trace; sharp 90º traces are a relic of the "tape-up" days of PCB layout. The rounding effects of CAD programs, however, still do not maintain constant width as the trace rounds the corner. Figure 7 shows progressively better techniques of rounding corners. Only the last example maintains constant trace width and minimizes reflections. Most CAD programs support these methods, but they can entail a little more work to master.

Trace-to-plane capacitors

PCB traces, being composed of foil, form capacitance with other traces that they cross on other layers. For two traces crossing each other on adjacent planes, this is seldom a problem. Coincident traces (those that occupy the same routing on different layers) form a long, skinny capacitor. Fortunately, these capacitances are usually small, only affecting high-frequency designs. It is important, however, to minimize capacitance at the inverting input to an op amp in high-speed designs. Otherwise, oscillation may occur.

Capacitance can be reduced by shortening the PCB traces and by reducing trace width. If oscillation still occurs, the resistors can be lowered a decade or two to damp it out. This will also result in lower circuit noise, at the penalty of increased power consumption.

Figure 7. Improving trace corners

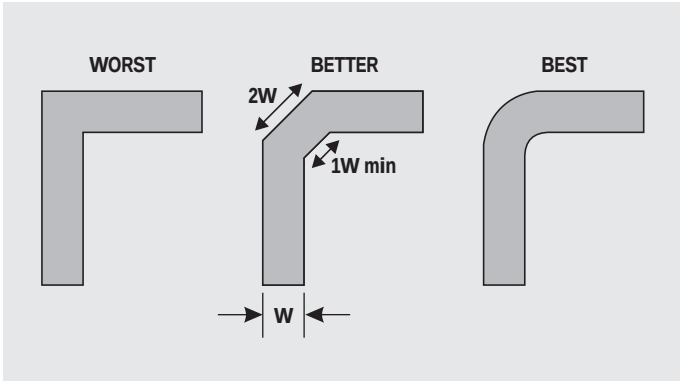

Trace-to-trace capacitors and inductors

PCB traces are not infinitely thin. They have some finite thickness, as defined by the "ounce" parameter of the copper clad foil. The higher the number of ounces, the thicker the copper. If two traces run side-by-side, then there will be capacitance and inductive

Continued on next page

Continued from previous page

coupling between them as shown in Figure 8. The formulas for these parasitic effects can be found in transmission line and/or microstrip references but are too complex for inclusion here.

Signal lines should not be routed parallel to each other, unless transmission line or microstrip effects are desired. Otherwise, a gap of at least three times the signal trace width should be maintained.

Figure 8. True behavior of parallel PCB traces

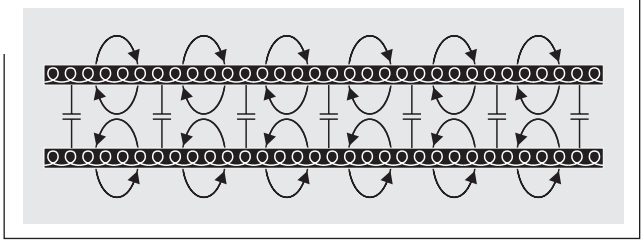

Inductive vias

Whenever routing restraints force vias, parasitic inductors are formed. This inductance is small but can become troublesome at high frequencies. It is best to avoid vias and route all signals on one layer of the board.

Flux residue resistance

An unclean board can affect analog circuit performance. If the circuit has very high resistances, even in the low megohms, pay special attention to cleaning. A finished assembly may be adversely affected by flux or cleansing residue. The electronics industry in the past few years has joined the rest of the world in becoming environmentally responsible. Hazardous chemicals are being removed from the manufacturing process, including flux that has to be cleaned with organic solvents. Water-soluble fluxes are becoming more common, but water itself easily can become contaminated with impurities. These impurities will lower the insulation characteristics of the PCB substrate. It is vitally important to use freshly distilled water every time a high-impedance circuit is cleaned. There are applications that may call for the older organic fluxes and solvents, such as very low-power, battery-powered equipment with resistors in the 10s of megohms range. Nothing can beat a good vapor defluxing machine for insuring that the board is clean.

Decoupling

Noise can propagate into analog circuitry through the power pins of the circuit as a whole and the op amp itself. If analog circuitry is located on the same board with digital circuitry, it is important to decouple the analog as well as the digital circuitry.

Table 1 is a rough guideline for the maximum useful frequencies of common capacitor types.

Table 1. Recommended maximum frequencies for capacitors

CAPACITOR TYPE	FREQUENCY
Aluminum electrolytic	100 kHz
Tantalum electrolytic	1 MHz
Mica	500 MHz
Ceramic	1 GHz

Tantalum electrolytic capacitors are useless for frequencies above 1 MHz. Effective high-frequency decoupling at higher frequencies demands a ceramic capacitor. Self-resonances of the capacitor must be known and avoided, or the capacitor may not help or may even make the problem worse. Figure 9 illustrates the typical self-resonance of two capacitors commonly used for bypassing —10-µF tantalum electrolytic and 0.01-µF ceramic.

Although these resonances are considered typical values, the characteristics of actual capacitors can vary from manufacturer to manufacturer and grade of part to grade of part. The important thing is to make sure that the self-resonance of the capacitor occurs at a frequency above the range of the noise you are trying to reject. Otherwise, the capacitor is inductive.

Don't assume that a single 0.1-µF capacitor will decouple all frequencies. Smaller capacitors may work better than larger ones at higher frequencies. When poor decoupling is suspected, try a smaller capacitor rather than a larger one.

A decoupling capacitor must be included on every op amp IC package. There are one, two, or four op amps per package. The value of the capacitor must be picked carefully to reject the type of noise present in the circuit.

In particularly troublesome cases, it may be necessary to add a series inductor into the power-supply line connecting

Figure 9. Typical self-resonance of 10-µF tantalum electrolytic and 0.01-µF ceramic capacitors

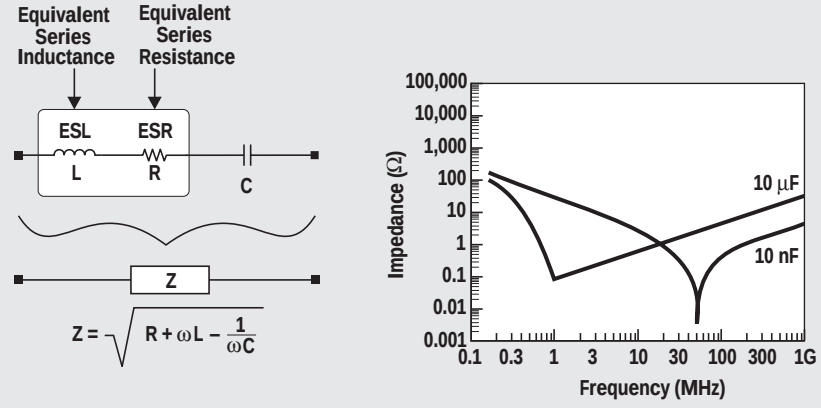

to the op amp. This inductor is in addition to the decoupling capacitors, which are the first line of defense. The inductor should be located before, not after, the capacitors.

Another technique that is lower in cost is to replace the series inductor with a small resistor in the 10-Ω to 100-Ω range. The resistor forms a low-pass filter with the decoupling capacitors. There is a penalty to pay for this technique; depending on the power consumption of the op amp, it will reduce the rail-to-rail voltage range. The resistor forms a voltage divider with the op amp as a resistive active component in the lower leg of the divider. Depending on the application, this may or may not be acceptable.

There is usually enough low-frequency ripple on the power supply at the board input to warrant a bulk decoupling capacitor at the power input. This capacitor is used primarily to reject low-frequency signals, so an aluminum or tantalum capacitor is acceptable. An additional ceramic cap at the power input will decouple any stray high-frequency switching noise that may be coupled from other boards.

Input and output isolation

Many noise problems are the result of noise being conducted into the circuit through its input and output pins. Due to the high-frequency limitations of passive components, the response of the circuit to high-frequency noise may be quite unpredictable.

In situations where conducted noise is substantially different in frequency from the normal operating range of the circuit, the solution may be as simple as a passive RC low-pass filter that rejects RF frequencies while having negligible effect at audio frequencies. A good example is RF noise being conducted into an audio op amp circuit. Be careful, though. A low-pass filter loses its characteristics at 100 to 1000 times f_{3db}. More stages may be required to cover different frequency ranges. If this is the case, make the highest-frequency filter the one nearest to the source of noise. Inductors or ferrite beads also can be used in a noise rejection filter network to eliminate conducted noise. Ferrite beads are inductive up to a specified frequency and then become resistive.

The effect of radiated energy coupling into an analog circuit can be so bad that the only solution to the problem may be to shield the circuit completely from radiated energy. This shield is called a “Faraday Cage” and must be carefully designed not to allow frequencies that are causing the problem to enter the circuit. This means that the shield must have no holes or slots larger that 1/20 the wavelength of the offending frequency. This can be a demanding requirement. It is a good idea to design a PCB from the beginning to have enough room to add a metal shield if it becomes necessary. If a shield is used, frequently the problem will be severe enough that ferrite beads also will be required on all connections to the circuit.

Packages

Op amps commonly are packaged one, two, or four per package. Single op amps often contain additional inputs for features such as offset nulling. Op amps packaged two and four per package offer only inverting and non-inverting inputs and the output. If the additional features are important, the only package choice is single. Be aware, though, that the offset-nulling pins on a single op amp package can act as secondary inputs and must be treated carefully. Consult the data sheet on the particular device you are using.

Figure 10. Dual op amps reduce long traces

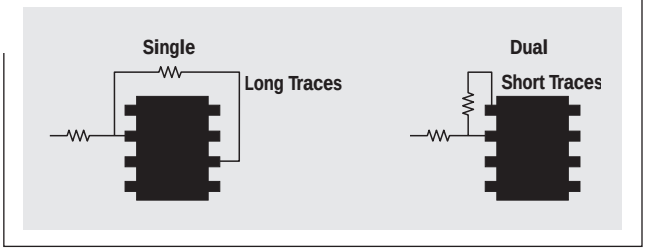

The single op amp package places the output on the opposite side from the inputs. This can be a disadvantage at high speeds because it forces longer PCB traces. It may make more sense to use a dual op amp package, even if the second op amp is unused. Figure 10 illustrates trace length for an inverting op amp stage.

Unused sections

In many op amp designs, one or more op amps may be unused. If this is the case, the unused section must be terminated properly. Improper termination can result in greater power consumption, more heat, and more noise in op amps on the same physical IC. If the unused section of the op amp is connected as shown under “Better” in Figure 11, it will be easier to use for design changes.

Figure 11. Proper termination of unused op amp sections

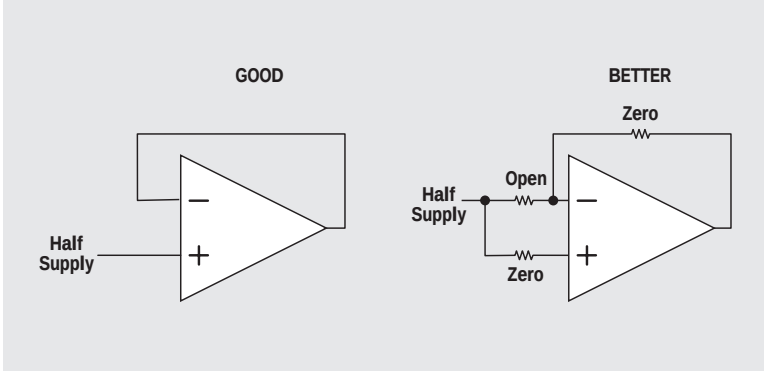

Related Web sites

www.ti.com/sc/docs/apps/analog/amplifiers.html

www.ti.com/sc/amplifiers

Reducing PCB design costs: From schematic capture to PCB layout

By Tom Hendrick
Engineering Technician, Advanced Analog Products

Modern EDA software packages provide a host of tools that allow designers to draw schematics in order to produce printed circuit boards. Understanding the basic requirements of these tools and the interactions between them can help to reduce the cost of PCB designs.

The basics

Schematic capture packages contain various tools and features that make the process of entering a schematic and documenting a circuit easy on the designer. One common feature among popular schematic capture packages is an array of libraries. Another is the ability to output netlists that are compatible with various simulation and PCB layout packages. The schematic capture process creates a database of symbolized parts and a netlist describing the connections between the symbols.

PCB layout packages have their own suite of tools and features designed to streamline the creation, verification, and documentation of a physical printed circuit board. All the board designer needs to do is define the outline of the board, add footprints from a decal library, import the netlist, and route the connections. Netlist comparisons will verify that the board matches the schematic. Online error checking warns of open- or short-circuit conditions. Design rules can be set up to check things such as matched net lengths, routing stubs, and parallelism.

Sounds simple, doesn't it? Draw a schematic, output the netlist, and sit back to wait for your prototype assemblies. Simple, that is, until the call from the assembly shop tells you that a part does not match its footprint, or you spend hours tracking down a diode or capacitor that was laid out backwards. An understanding of the process that takes place when exporting a schematic netlist is necessary to ensure your board layout will be correct.

Figure 1. Schematic library symbols

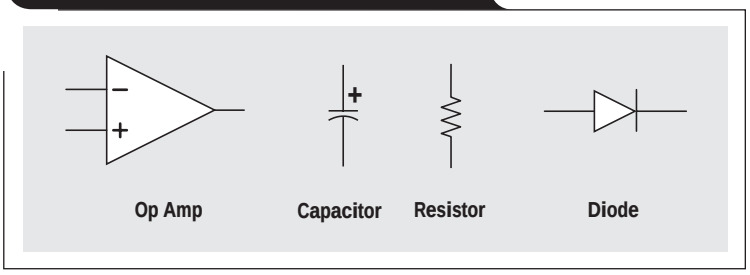

Figure 2. Component electrical pin numbers

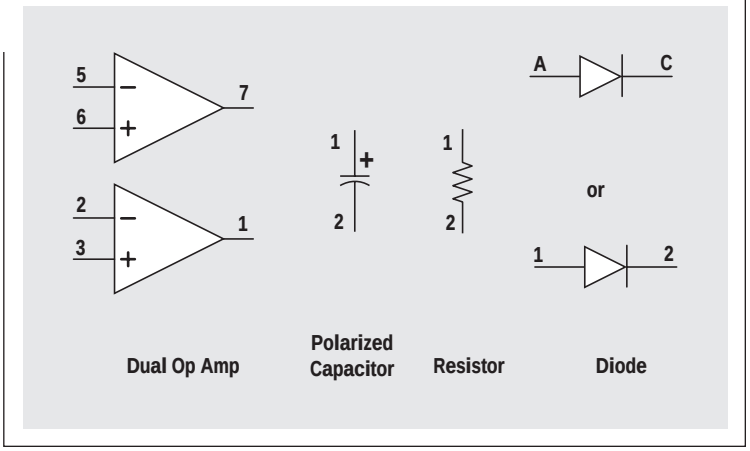

Schematic libraries

Schematic libraries are based on familiar symbols such as those shown in Figure 1.

The symbol libraries are generic by nature. They are intended as building blocks for various part types. Schematically speaking, the same resistor symbol can be used for a 1/8-watt surface mount or a 100-watt chassis-mount device. The op amp symbol could be a single device in a 5-pin SOT package or a dual device in an 8-pin DIP. To create a valid schematic part, these symbols must, at a minimum, have electrical pin numbers assigned and be saved in a part type library.

Quite often, this is the extent of the part type libraries supplied with the schematic capture tool (see Figure 2). Enough information is given to allow the designer to create a schematic, but information required to create a printed circuit board is missing. As with the symbol library, the vendor-supplied part type libraries are intended as a starting point for a customized user library.

Customizing a vendor-supplied library requires that the user assign attributes detailing the specific part to be used in the design. A generic vendor-supplied part type "resistor," for example, could be copied to a customized user library, be given a footprint attribute "1206," and then be renamed to part type "RES_1206_SMT." When this new part is called from the user library, it would always require that a 1206-size surface-mount footprint be used to represent this device on the printed circuit board.

PCB libraries

The PCB libraries are based on the footprints designated by industry standards. They usually include common surface-mount and through-hole devices. To create a footprint, or decal, pads are arranged in a specific pattern and designated with an electrical pin number. Figure 3 is

an example of some common footprints shipped with most PCB design tools.

As with the schematic symbol library, one PCB decal can be associated with many different part types. The layout database relies on the pin numbers defined in the schematic library for mapping to a footprint in the PCB library. This process allows two completely different 8-pin devices to use the same 8-pin PCB decal.

Understanding the importance of this problem becomes even more obvious when you consider parts such as diodes or electrolytic capacitors. A generic schematic part called "DIODE," for instance, could be assigned a PCB decal called "D0-41." If the diode is designated in the schematic library with electrical pin 1 as the anode and electrical pin 2 as the cathode, the associated D0-41 footprint must also have the same designations (see Figure 4). Having the pins reversed on the PCB footprint can cause serious problems that can be difficult to debug on an assembled board. Error checking will not catch this type of problem.

Figure 3. Example of common PCB footprints

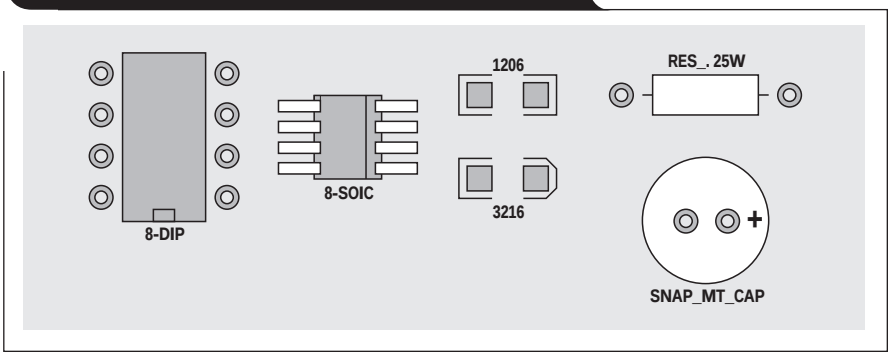

Figure 4. Diode designation

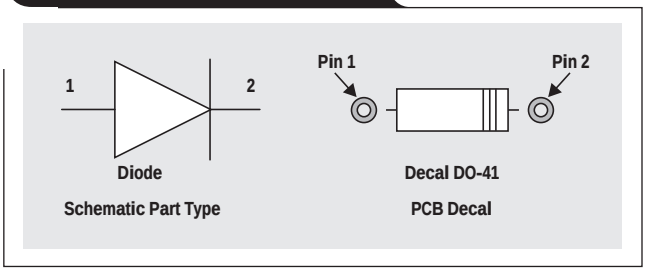

The netlist file

The netlist file is where the schematic and PCB databases come together. The output from a schematic database is formatted to meet the needs of a specific PCB layout package and contains several specialized sections.

The first section usually describes the PCB decals that are needed in the design. If a decal does not exist in the PCB library, the desired part cannot be placed on the board. All schematic part types must be assigned a valid PCB decal before the schematic netlist can be imported to the PCB design database. Most PCB design packages will immediately flag an error and stop processing the netlist file if a decal is missing.

The next section of the netlist file contains the connection, or "rubber band," information. If a decal cannot be found, there can be no connections to route. Other sections may include PCB routing directives, simulation directives, or detailed information for power and ground connections.

Save time and avoid errors

Avoid using different platforms to produce schematic and PCB designs. Most EDA vendors can provide bundled schematic capture and PCB layout packages that are designed to work together. These integrated packages provide a common shell that allows for easier verification of the schematic and board files, and makes processing engineering changes relatively painless. Integrated packages usually allow "back-annotation" as well, so that changes made in a PCB design can be reflected back to the schematic.

The use of a common library can help prevent problems. Develop a system where one regularly updated library serves the needs of all persons involved in design activities. Multiple people creating the same devices in individual libraries will only lead to errors and confusion at layout. If a new part needs to be constructed, build it in a central library so everyone can use it. Review all newly created parts and their associated PCB footprints for accuracy.

PCB layout tools will normally present an error or warning if the schematic part references a PCB footprint with fewer electrical pins than described in the schematic netlist. For example, if a 14-pin device were assigned to an 8-pin decal, there would be an error mapping pins 9 through 14 described in the schematic. If, on the other hand, a 14-pin decal is assigned to an 8-pin device, error flags are not necessarily raised. A PCB decal can have more pins than electrically represented in a schematic; mounting or alignment holes, for instance, are part of a decal. This type of error often goes unnoticed as well. It is usually found after the PCB is fabricated, when someone tries to mount a 14-pin device in a 16-pin decal.

Many companies use outside sources to design printed circuit boards. If an external design house provides your printed circuit board layouts, make sure your engineers are familiar with the tools they use. Sending out a netlist file in the correct format will save time and money. Share your libraries with your vendor so they can save time and avoid having to re-create parts. Create your schematic libraries so that they reference the PCB footprint libraries of your design house.

TI Worldwide Technical Support

Internet

TI Semiconductor Home Page
www.ti.com/sc

TI Distributors
www.ti.com/sc/docs/general/distrib.htm

Product Information Centers

Americas

Phone	+1(972) 644-5580
Fax	+1(214) 480-7800
Internet	www.ti.com/sc/ampic

Europe, Middle East, and Africa

Phone	
Belgium (English)	+32 (0) 27 45 55 32
France	+33 (0) 1 30 70 11 64
Germany	+49 (0) 8161 80 33 11
Israel (English)	1800 949 0107
Italy	800 79 11 37
Netherlands (English)	+31 (0) 546 87 95 45
Spain	+34 902 35 40 28
Sweden (English)	+46 (0) 8587 555 22
United Kingdom	+44 (0) 1604 66 33 99
Fax	+44 (0) 1604 66 33 34
Email	epic@ti.com
Internet	www.ti.com/sc/epic

Japan

Phone	
International	+81-3-3344-5311
Domestic	0120-81-0026
Fax	
International	+81-3-3344-5317
Domestic	0120-81-0036
Internet	
International	www.ti.com/sc/jpic
Domestic	www.tij.co.jp/pic

Asia

Phone		
International	+886-2-23786800	
Domestic	Local Access Code	TI Number
Australia	1-800-881-011	-800-800-1450
China	00-800-8800-6800	-
Hong Kong	800-96-1111	-800-800-1450
India	000-117	-800-800-1450
Indonesia	001-801-10	-800-800-1450
Korea	001-800-8800-6800	-
Malaysia	1-800-800-011	-800-800-1450
New Zealand	000-911	-800-800-1450
Philippines	105-11	-800-800-1450
Singapore	800-0111-111	-800-800-1450
Taiwan	080-006800	-
Thailand	0019-991-1111	-800-800-1450
Fax	886-2-2378-6808	
Email	tiasia@ti.com	
Internet	www.ti.com/sc/apic	

A050200