High-Performance Analog Products

Analog Applications Journal

First Quarter, 2006

IMPORTANT NOTICE

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

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

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

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

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

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

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

Products	
Amplifiers	amplifier.ti.com
Data Converters	dataconverter.ti.com
DSP	dsp.ti.com
Interface	interface.ti.com
Logic	logic.ti.com
Power Management	power.ti.com
Microcontrollers	microcontroller.ti.com

Applications	
Audio	www.ti.com/audio
Automotive	www.ti.com/automotive
Broadband	www.ti.com/broadband
Digital control	www.ti.com/digitalcontrol
Military	www.ti.com/military
Optical Networking	www.ti.com/opticalnetwork
Security	www.ti.com/security
Telephony	www.ti.com/telephony
Video & Imaging	www.ti.com/video
Wireless	www.ti.com/wireless

Mailing Address: Texas Instruments
Post Office Box 655303
Dallas, Texas 75265

Contents

To view past issues of the
***Analog Applications Journal*, visit the Web site**
www.ti.com/aaj

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 categories:

- Data Acquisition
- Power Management
- Amplifiers: Op Amps

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.

Understanding and comparing datasheets for high-speed ADCs

By Eduardo Bartolome (Email: e-bartolome1@ti.com)
High-Speed ADC Systems and Applications Manager

Introduction

As with any other aspect of a product, datasheets can always be improved, and manufacturers work hard to clarify them.[1] Nevertheless, there is already a legacy of many products/datasheets in the market, and even for those recently released or under release, different criteria may have been selected depending on different factors. Even given the existence of a few published standards such as Reference 2, convergence is still far away.

The purpose of this article is to highlight the differences between the criteria for writing datasheets applied by different manufacturers or even by the same manufacturer of different high-speed analog-to-digital converters (ADCs). Table 1 is a quick checklist that can be used to select the right device.

dB versus dBc versus dBFS

One of the main criteria for evaluating high-speed-ADC performance shown on a datasheet is the dynamic or ac set of specifications like SNR, SINAD, THD, and SFDR. As an example, let's consider the SINAD (also called SNDR or SNRD), which is the ratio of signal power to the power of all other spectral components (distortion and noise) combined. Usually, such a ratio is transformed into logarithmic scale and expressed in dB. In this particular case,

$$\text{SINAD} = 10 \times \log\left(\frac{P_S}{P_{N+D}}\right).$$

One of the biggest discrepancies between datasheets arises in the selection of units, specifically in the criteria used to select the power of the signal. Without going into the details that lead to this disparity, let's note that the power of the signal could be simply the power of the signal being used on the measurement,[3] or it could be the extrapolated power of the signal as if it was full-scale. To avoid confusion, Texas Instruments (TI) indicates the units in the first case as dBc (dB to the carrier), and in the second case as dB in older datasheets or as dBFS (dB to full scale) in the most recent datasheets. Although there is no major issue between the two and one can be obtained from the other as will be explained later, the confusion arises when the manufacturer specifies only "dB." In the datasheets of two TI competitors, "dB" is equivalent to dBc, while in the datasheet of a third competitor, "dB" is equivalent to dBFS.

What impact does this have on the final comparison? Let's imagine that two converters from different manufacturers have been measured with an input amplitude of –1 dBFS. This means that during the test, the input was a sinusoidal tone 1 dB below full scale or, in other words, had an amplitude (A) about 90% of the ADC's full dynamic range. Let's also assume that both converters have the same performance and, as such, the noise plus distortion power is the same. In the first case ("dB" is equivalent to dBc), the manufacturer computed the SINAD_1 as the power of the input (A) divided by the power of noise plus distortion (N+D). In the second case ("dB" is equivalent to dBFS),

Table 1. Quick checklist for selecting high-speed ADCs

SPECIFICATION	CONSIDERATION
Power consumption	Make sure that output supply is included and is specified under the same conditions in which you are going to use the device.
Input clock amplifier	If you are going to sample high-input frequencies, look for differential clock inputs and an input clock amplifier.
SNR, SINAD, SFDR, THD	Be careful with dB units. They may mean dBc or dBFS.
Output timing	Make sure output timing is properly specified and has enough window to capture the data. Some devices offer an output clock with a better timing window for easier interface.
Input bandwidth	Device performance is not necessarily guaranteed up to the frequencies specified.
General specifications	• Since the manufacturer will ensure only minimum values, look for those rather than typical values. • Check for specifications under the same conditions in which you are going to use the device. • Look for plots showing device robustness, like performance variation over different temperatures or supply voltages. • Verify the performance versus the input amplitude graph if required.

the manufacturer extrapolated the power of the signal to the full scale, which in this case resulted in an increase of the final number by 1 dB (as the input was –1 dBFS). So, $SINAD_2 = SINAD_1 + 1$. The first manufacturer's device looks 1 dB worse but is actually as good as the one from the second manufacturer. The dBc value can be extracted from the dBFS value by adding the input amplitude (in dBFS):

Specification (in dBc) = Specification (in dBFS) + A_{IN} (in dBFS).

Another point to take into account when we compare two devices whose specifications are expressed in dBc is the input amplitude, which also changes from device to device. Obviously, as the input amplitude decreases, the signal level (in dBc) will most likely become smaller (as the signal power becomes smaller while some components of the noise floor stay the same). So, for the comparison to be meaningful, both input amplitudes should be the same. A good extrapolation can be done by just adding the difference between both input amplitudes to the specification obtained with the smaller input amplitude.

Observe that as most of the converters are specified with signals close to the full scale, the difference in the final signal level, depending on the units being used, will usually not be bigger than 2 dB, although that can actually represent a considerable difference in some applications, affect the yields of the final design in production, or even force a whole redesign.

The final question is how to distinguish whether the dB in the ADC datasheet is dBc or dBFS. One of the easiest ways is to ask the vendor directly; but, if that is not possible, the designer can look at the typical-performance graphs section of the datasheet. There vendors will usually have different graphs in dBc and dBFS whose numbers can be compared to the ones specified in the datasheet tables.

Power specifications

Power is another point where datasheets usually show discrepancies. TI has traditionally listed on the first page of its datasheets the total power of the converter under typical operating conditions. These conditions are at full sample rate with an input different than the low frequency or dc, and, on the digital side, with a 10-pF load on each pin. However, TI has started separating core (analog, AV_{DD}) power and digital (output, DV_{DD}) supply consumptions. This is mainly because almost every other vendor specifies only the analog consumption and does not include the interface power. Some vendors list the interface power inside the datasheet; but users should pay attention to the conditions of such a measurement, as input frequency, output load per pin, and output voltage affect that value significantly.

Also, special attention should be paid to special modes that the ADC may have. Although such is not the case with TI datasheets, some vendors' converters have a mode A under which the power is measured and listed in the datasheet, while all the other dynamic specifications are actually given under a mode B. Or, for instance, for converters that accept wide digital supply voltages, the datasheet gives the specification of the timing under the highest DV_{DD} but provides the power consumption numbers under the lowest DV_{DD}.

Finally, pay attention to what the converter offers for that power. Some converters may save power at the expense of missing internal blocks such as the reference or clock amplifier that could be needed for things like accepting differential clocking (for lower noise/jitter), small swing clock levels like PECL or LVDS, or squaring a filtered clock (sinusoidal). More on this topic follows.

Clocking

Clocking is a fundamental care-about to obtain the best ADC performance[4]; and, although all ADCs have a clock input, some are easier to use than others. Specific concerns are clock jitter, duty cycle, and required clock levels, mainly to obtain good performance numbers when high input frequencies are sampled.

To obtain low jitter, users usually bandpass filter the clock. This will produce a 50% duty cycle, which is usually close to the best condition for many ADCs. Nevertheless, the clock amplitude will be affected due to the insertion loss of the filter; and the shape will become sinusoidal, not square. To be able to accept this high-quality clock, TI and some other manufacturers include a clock amplifier at the input of the ADC. Its functions, among others, is to square the sinusoidal clock and provide the necessary gain to clock the internal circuits. The clock amplifier will, moreover, present a differential interface that will make it less susceptible to noise coupling in the clock lines, reducing jitter. Of course, this will come at the expense of some increased ADC power consumption.

Nevertheless, some ADCs may not be so friendly and will require a square CMOS-level clock. The input on these types of ADCs will be single-ended, with zero rejection of external noise sources coupling to the clock path. Most of these ADCs are targeted to the sampling of low (below 50 MHz or so) input frequencies and yield a good performance there. Medical ultrasound would be a typical application for these ADCs. Nevertheless, users targeting high (greater than high 60s) SNR at higher input frequencies (like on communication applications) should know that they will need to provide external components to square the clock, effectively increasing the power and the board area. Even with those extra components, there is nothing that the user can do with respect to the single-ended interface, and the final design most likely will not obtain the optimized performance that ADCs with an internal clock amplifier provide.

In the effort to provide consistent datasheets, TI actually uses all the plots and performance numbers on the same datasheet that share the same clock conditions. For devices with a differential input clock, usually a sinusoidal is used, which is usually not the best condition for the ADC (due to the limited slew rate of the clock edges). To cover all possible applications of the converter, TI started

introducing 3D contour plots (Figure 1) in its datasheets, which allow the customer to find the typical performance under a given set of input and sampling frequencies. As just noted, all the conditions of the input clock are kept constant during the experiment except, obviously, when its sampling frequency is changed. This means that if a sinusoidal clock is used, reducing the sampling frequency will make the edges of the clock become slower, increasing the effective jitter. This is a common phenomenon in every ADC,[4] but TI's ADCs are designed to minimize it as much as possible. Although this is a worst-case situation and there are techniques to minimize jitter, it would not be fair to change the clock conditions during the experiment; and, as such, they are kept constant and show the SNR degradation due to the increased jitter. Users should be aware that without the clock amplifier, the performance degradation would actually be much worse. They should also be aware that if a very low-jitter square clock could be provided, the performance of the ADC would actually be much better.

Figure 1. SNR vs. input and sampling frequencies[5]

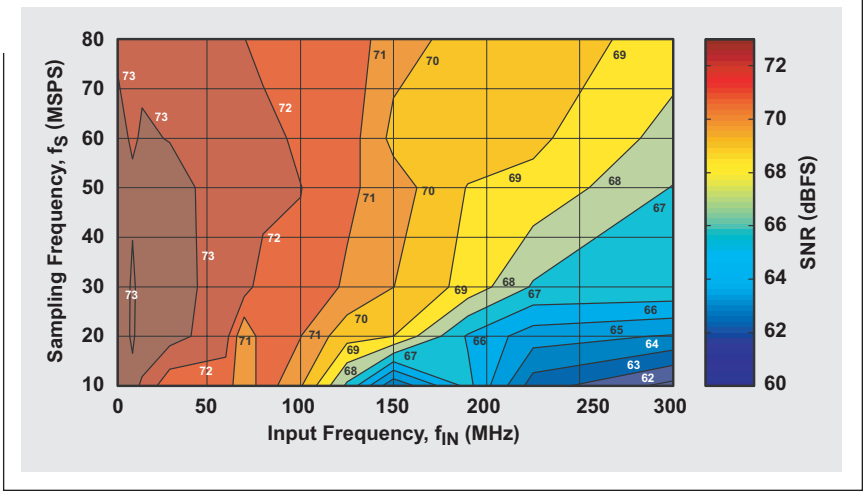

Output timing

To be able to capture the output data with the FPGA, ASIC, DDC, or any other logic device following the ADC, users need to know the window where the output data is stable. Nevertheless, this is a point where most of the manufacturers struggle to provide a consistent and complete set of datasheet limits. This is because the final test solution being used in production is affected by several factors such as precision of the automated test equipment, no direct access to the outputs (data being buffered), difficulty setting the same conditions (like digital load) as the datasheet, etc. To bypass these limitations, TI set these numbers by design and characterization—i.e., statistically—which usually forces us to set wider guardbands, as the devices are not going to be tested in production. The same type of limitations apply to other manufacturers, many times resulting in incomplete or poorly specified datasheets.

Designers should specifically question devices that do not have any guarantees at all; do have guarantees but under unrealistic conditions (like 0-pF load); do not specify all the parameters necessary for the capture (for instance, set-up time is given but not hold time); do not explain what V_{OH} and V_{OL} levels are used with the specifications (for example, giving information from 50% to 50%, which is very cumbersome to extrapolate to V_{IH}/V_{IL} logic levels); or do not specify a parameter over the full operating temperature range.

Also, to improve the window for data capture, TI and other manufacturers provide an output clock that has better tracking with the output data than the input clock. Using the output clock will alleviate the timing constraints on the application.

Finally, notice that set-up and hold times are defined as the counterparts of the set-up and hold times defined for latches. In a latch, set-up time indicates how early the data has to be ready at the input of the latch before the clock edge is provided to latch it. The longer this time is, the harder it is to use that latch. On the ADC, the set-up time indicates how early the data is stable before the input or output ADC clock edge. The longer the set-up time is, the easier it is to capture the data. The same applies to the hold times.

Process gain

When comparing two SNR numbers, users should take into account the sampling rate of the ADC. The SNR is computed by integrating the total noise floor all the way to Nyquist. Nevertheless, the user's signal will occupy only a certain bandwidth; and only the noise on that bandwidth will affect the signal, as the rest can be removed with a digital filter. For the same SNR, the ADC with a higher sample rate will have a lower noise floor. For example, a 200-kHz-bandwidth sigma-delta with 90-dBFS SNR would look better on paper than the ADS5424, a 14-bit ADC running at 100 MSPS with a 75-dBFS SNR on its datasheet. If, however, after using the ADS5424 to sample the 200-kHz bandwidth (clearly oversampling at 100 MSPS), we use a digital filter to remove the out-of-band noise (from 200 kHz all the way up to 50 MHz), the equivalent SNR of the ADS5424 on that bandwidth will be

$$\begin{aligned} SNR_{200\ kHz} &= 75 + 10 \times \log10(50 \times 10^6/200 \times 10^3) \\ &= 99\ dBFS >> 90\ dBFS. \end{aligned}$$

The SNR of the ADS5424 would clearly be much better than that of the sigma-delta (assuming, for the sake of this example, that the noise is evenly distributed over Nyquist; i.e., there is no significant flicker contribution). The second term of the preceding equation is called process gain. As the oversampling ratio increases, for every time the user doubles the sampling rate for the same SNR, the noise

floor will decrease by 3 dB. In other words, the equivalent resolution on the bandwidth of interest will be increased by 0.5 bit.

Conditions and guaranteed minimums

A value/specification is meaningless if the conditions are not expressed. This is especially important when we look at advertisements, marketing materials, and selection tables, which tend to simplify the more complete information presented in a datasheet and, in some cases, do not mention the conditions of the measurement.

At the same time, typical values usually represent the average of a distribution. Nevertheless, users should look at the minimum values, especially if the device is going to operate over a certain range of conditions like varying temperature. Having a wide spread between the typical and minimum numbers should raise questions. Is the spread due to a limitation of the final test solution? If so, maybe the device will fit the application but with a certain risk that only the minimum values are guaranteed. The worst case would be if the limitation was due to the device itself. A process variation (from device to device) could point to a weak design. A good way to look for the robustness of the device and screen for these types of problems is to look at the datasheet graphs of typical performance variations versus supply voltage or temperature, similar to the example in Figure 2.

Just as with the power numbers, when looking at different parameters (for instance, SNR and SFDR) in the same datasheet, make sure they are all given under the same conditions. For instance, some devices have an SFDR mode where the SFDR improves at the expense of the SNR; or they have different input ranges, which can affect both the SFDR (better at the smaller input range) and the SNR (better at the bigger input range).

Finally, be aware that most of the specifications are given close to full-scale range. Nevertheless, the SFDR (in dBFS) may get better or worse as the input amplitude is decreased. Manufacturers cannot screen many conditions on the final test, as that would increase the test time and thereby the cost; but they usually include typical graphs that show the effects of various conditions.

Figure 2. ADS5424 SFDR vs. analog supply and temperature range[6]

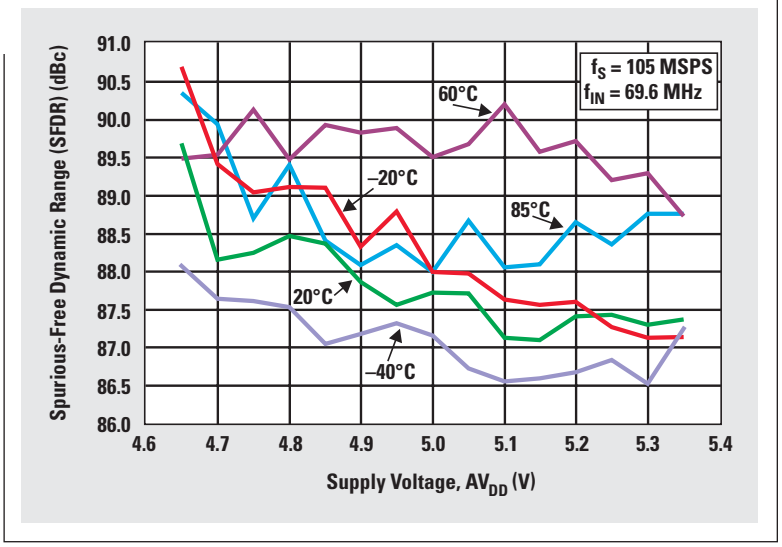

Input bandwidth

Usually the input bandwidth is representative of the flatness of the ADC response versus input frequency. It is not an indication that the device will hold the performance up to those input frequencies. Users should look for graphs that verify performance; and, if none exist, they should request support from the vendor or evaluate the device themselves.

Conclusion

A detailed summary of the main differences between criteria that manufacturers use to write datasheets for high-speed ADCs has been presented. Early consideration of these differences will help designers avoid surprises further down the road.

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. Joshua Israelsohn, "High-speed ADCs: preventing front-end collisions," *EDN Magazine*, May 13, 2004.	—
2. IEEE Standard 1241, "IEEE Standard for Terminology and Test Methods for Analog-to-Digital Converters."	—
3. Hui-Qing Liu, "14-bit, 125-MSPS ADS5500 evaluation"	slyt074
4. Eduardo Bartolome et al., "Clocking high-speed data converters"	slyt075
5. "14-bit, 125-MSPS Analog-to-Digital Converter," ADS5500 Datasheet, p. 17	sbas303
6. "14-Bit, 105-MSPS Analog-to-Digital Converter," ADS5424 Datasheet, p. 14	slws157

Related Web sites

dataconverter.ti.com
www.ti.com/sc/device/ADS5424
www.ti.com/sc/device/ADS5500

Powering today's multi-rail FPGAs and DSPs, Part 1

By Jeff Falin (Email: j-falin1@ti.com)
Applications, Portable Power

Introduction

Most electronics have one or more digital processing ICs, such as FPGAs or DSPs, that require multiple power-supply rails. There are various options to consider and potential pitfalls to avoid in powering these digital ICs. This article, Part 1 of a two-part series, provides recommendations and guidance for developing a power solution for multi-rail applications where the input power supply voltage is assumed to be equal to or greater than the system rail voltage (e.g., 12, 5, or 3.3 V). Part 2, which will appear in a future issue of *Analog Applications Journal*, will focus on how to choose between the types of dc/dc converters and how to design them to meet dc accuracy, start-up, and transient requirements.

Application-specific requirements

Application-specific requirements drive the overall dc/dc power solution. While most system designers desire a simple power solution with low component cost, they must also consider the type of circuitry being powered as well as the differential between the input power supply and each power rail. FPGA and DSP core and I/O voltage rails already have switching noise riding on the rail due to the millions of internal transistors that are switching on and off. So, in general, these "digital" rails can be powered from switching power supplies without concern for the switching noise. Conversely, rails that power audio circuitry, transceiver circuitry, clock signals, phase-lock loops, or other noise-sensitive circuitry—termed "analog" rails—should be powered from a linear regulator or low-noise, fixed-frequency PWM converter. These rails may even require additional filtering on the power rail as specified by the analog IC vendor. In addition, noise-sensitive circuitry may be affected by noise emitted from switching regulators, so recommended board layouts for switching regulators should be followed and shielded inductors used. Restrictions on the range of switching frequency and/or synchronization of all switching regulators to one switching frequency may be necessary to make filtering easier. Once the type of circuitry being powered has been reviewed, the differential between the input power supply voltage and each power rail voltage—and therefore the power dissipation that each rail's converter must withstand—must be considered. An easy way to do this is to prepare a power budget.

The power budget

Shrinking process nodes have resulted in core rails dropping from 2.5 to 1.2 V or even below 1.0 V. However, wall-brick voltages and many bus voltages used as the input supply to the point-of-load (POL) dc/dc converters that provide the core voltages have remained at 12, 5, and 3.3 V. The dc/dc converters providing such voltages must be able to handle the power dissipation. Also, as the efficiency of the individual POL converters drops, the output power of the input power supply increases. Although the static or quiescent current for a digital IC is typically provided in the datasheet, estimating the maximum current for a specific FPGA configuration or software program (sometimes referred to as "dynamic current") is often difficult to do. Fortunately, all FPGA and DSP manufacturers provide either online power-consumption estimators or downloadable software that takes key, application-specific design criteria and provides at least a working estimate of the maximum current draw. Once the current, and therefore power, consumption for each rail is known, a power solution can be constructed. The block diagram in Figure 1 shows one of many possible power solutions for an application containing one FPGA and one DSP.

Figure 1. Simplified application with one FPGA and one DSP

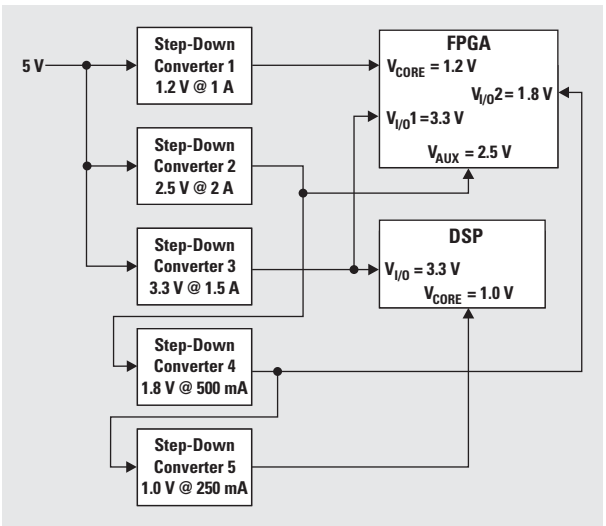

Note that instead of pulling all of the rail voltages from the 5-V input supply, the design uses the FPGA's required 2.5-V rail as a mid-level bus voltage to more efficiently provide the 1.8-V rail and then uses the 1.8-V rail to provide the 1.0-V rail. If all of the POL converters are assumed to be inexpensive linear regulators, then the power budget in Table 1 shows that the input supply must provide at least 4.5 A. Also, converters 1 and 2 must dissipate 3.8 and 5 W, respectively, which is difficult for a surface-mount-packaged linear regulator to do without additional air flow or an external heat sink.

The power budget in Table 2 assumes that converters 1, 2, and 3 are buck-switching converters. With higher efficiencies as well as reduced input current, switching regulators eliminate power dissipation concerns and allow the use of a lower-power, less expensive wall brick for the input supply.

Good power-supply design techniques

Designing a multi-rail-powered system without considering the power-up timing of each rail poses several dangers that can threaten immediate and long-term device reliability. Improper rail timing can magnify latent faults and possibly damage I/O ports in the processor or supporting system devices such as memory, logic, or data-converter ICs. A long-term threat comes from the possibility of breakdown in the ESD protection and well-isolation structures that internally separate the two power-supply rails. If one rail is active while another is inactive for extended periods—i.e., months—or for repeated shorter periods that accumulate over the lifetime of the device, damage may occur.

A latch-up condition may cause damage that is immediately noticeable, or it may affect reliability over a long period of time. Latch-up occurs when current forced through the substrate of a CMOS device triggers a self-sustained conduction path in back-to-back, parasitic bipolar transistors (as in an SCR). Current continues to flow until the device fails or the supply powers down. The trigger current may occur if a supply powers one device, enabling it to source or sink current into or out of the second device before the second device fully powers up. A system can also trigger a latch-up if it drives an input pin above or below the power-supply rails after both devices have powered up. Bidirectional I/O ports have historically been another source of system failures. When a processor's I/O port and that of a supporting peripheral, such as memory or a data converter, do not share the same supply, the potential for latch-up exists. Bus contention occurs when multiple devices simultaneously attempt to control a bidirectional bus during power-up, which can affect I/O reliability. Similar undesirable conditions where a "sneak" path from the rising input rail(s) to ground is temporarily created by transistors in unknown states can cause a digital device to pull large inrush currents, which may cause immediate damage or cumulative long-term reliability concerns. Recently, FPGA and DSP manufacturers have improved protection circuits to reduce the risk of latch-up, bus contention, and similar undesirable states.

As further protection against these problems, the following simple-to-implement power-supply design practices are recommended. First, ensure that any logic peripheral connected to the processor I/O bus is powered from the same supply rail that powers the processor's I/O section. Second, ensure that the individual converters providing the rails regulate within the specified tolerance over the entire input voltage and load range during a transient

Table 1. Power budget assuming that all linear regulators are used

POL CONVERTER	V_{OUT} RAIL (V)	ESTIMATED I_{OUT} (A)	P_{OUT} (W)	POL CONVERTER V_{IN} (V)	ESTIMATED EFFICIENCY (%)	$P_{IN} = P_{OUT}/EFFICIENCY$ (W)	I_{IN} REQUIRED $= P_{IN}/V_{IN}$ (A)	POWER DISSIPATED (W)
5 Linear	1	0.25	0.25	1.8	56	0.45	0.25	0.20
4 Linear	1.8	0.5	0.9	2.5	72	1.25	0.50	0.35
3 Linear	3.3	1.5	4.95	5	66	7.50	1.50	2.55
2 Linear	2.5	2	5	5	50	10.00	2.00	5.00
1 Linear	1.2	1	1.2	5	24	5.00	1.00	3.80
TOTAL						**22.50**	**4.50**	

Table 2. Power budget assuming that switchers and linear regulators are used

POL CONVERTER	V_{OUT} RAIL (V)	ESTIMATED I_{OUT} (A)	P_{OUT} (W)	POL CONVERTER V_{IN} (V)	ESTIMATED EFFICIENCY (%)	$P_{IN} = P_{OUT}/EFFICIENCY$ (W)	I_{IN} REQUIRED $= P_{IN}/V_{IN}$ (A)	POWER DISSIPATED (W)
5 Linear	1	0.25	0.25	1.8	56	0.45	0.25	0.20
4 Linear	1.8	0.5	0.9	2.5	72	1.25	0.50	0.35
3 Switching	3.3	1.5	4.95	5	93	5.32	1.06	0.37
2 Switching	2.5	2	5	5	90	5.56	1.11	0.56
1 Switching	1.2	1	1.2	5	80	1.50	0.30	0.30
TOTAL						**12.38**	**2.47**	

event. This is explained more fully in Part 2 of this series. Third, ensure that all of the rails power up monotonically within a relatively short time (typically ≤ 100 ms) or the same time as each other so that the time of unexpected differential between the rails during startup is minimized.

Implementing sequencing

Sequencing refers not only to the order in which voltage rails power up and down but also to their timing and voltage-differential relationships. FPGA and DSP manufacturers rarely require a specific sequencing order for their power rails, but when they do, it should always be followed. A recommended sequencing order usually means that the manufacturer has performed at least a limited set of successful power-up tests with that device.

There are three methods for controlling power-up sequencing: simultaneous, ratiometric, and sequential. Figure 2 shows an example of simultaneous sequencing in which both rails ramp simultaneously, with the lower rail stopping at its regulation point. The rails initially rise at the same dv/dt, and the time that each rail is outside of its regulation tolerance is minimized. Therefore, simultaneous sequencing is generally considered to be the ideal sequencing method because it prevents latch-up, bus contention, and undesirable transistor states. However, it is also the most difficult to implement without special circuitry that allows interaction between each converter.

Ratiometric sequencing is shown in Figure 3. In ratiometric sequencing, the rails rise at different dv/dt rates, with the higher rail having a higher slew rate, so that they each arrive at their regulation point at the same time. Also, note that the maximum voltage differential occurs at the point of regulation. This sequencing method can be easily implemented if the dc/dc converters have an externally controllable soft start (discussed later).

Figure 4 shows the easiest and simplest method for soft-start sequential sequencing. By simply tying the output voltage (or available power good signal) of the first converter to the second converter and so on, sequential sequencing can be implemented. Maximum differential between the voltage rails can occur with this method; but as long as all of the rail voltages rise quickly enough relative to each other, the risk of immediate damage or long-term reliability problems is negligible. In addition, many of the recommended sequencing requirements for FPGAs and DSPs are sequential, since this is the simplest method around which to build a test setup.

Power-down sequencing is difficult to manage because the rail capacitance and resistive loading determine the voltage profile of each rail during power down. In other words, even if rail 1 is disabled before rail 2, rail 1 may have a larger capacitive load and smaller resistive load than rail 2; so rail 1 will discharge after rail 2. In addition, power down can occur as a result of several scenarios, including situations in which the loading on either rail changes during power down and is different from one power-down event to another. Additional circuitry is necessary that either discharges the rail voltages with a

Figure 2. Simultaneous sequencing

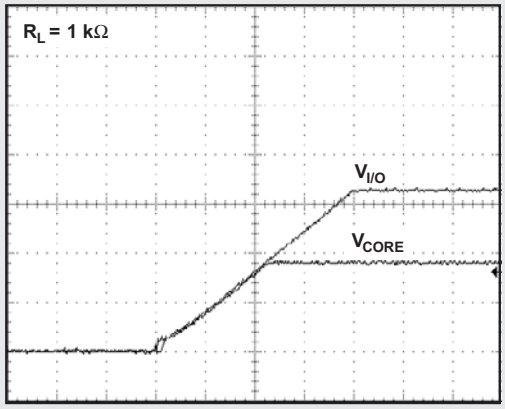

Figure 3. Ratiometric sequencing

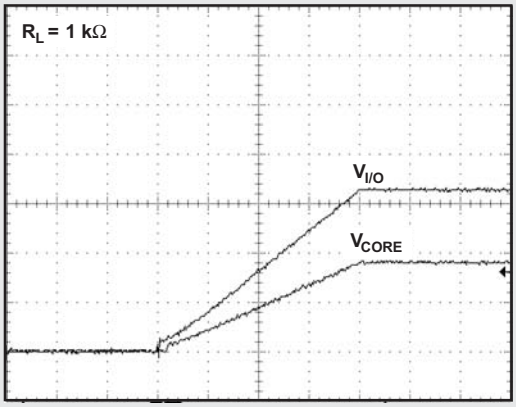

Figure 4. Sequential sequencing

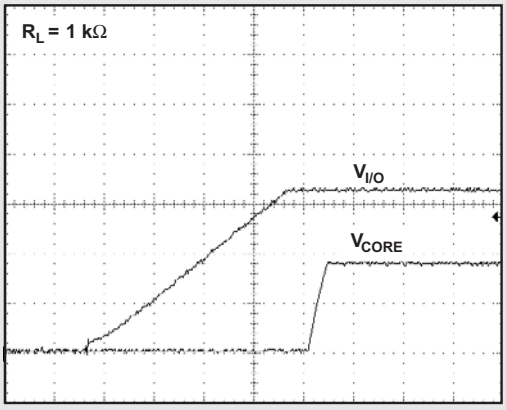

known load each time or actively monitors all rails and forces them to track each other (e.g., simultaneous sequencing). The smooth, monotonic rise of each individual rail, as Figures 2–4 show, may also be a requirement.

Implementing a controlled monotonic rise of the power rail

A monotonic rise at startup is shown in the top graph of Figure 5. Many FPGAs use an internal supply voltage supervisor (SVS) on each power rail, so a monotonic rise is necessary for the SVS circuit to successfully power up the rest of the IC. Many DSPs require an externally provided SVS or RESET signal that indicates when all the power rails are up, so monotonic rise is less important.

Figure 5. Start-up voltage and current

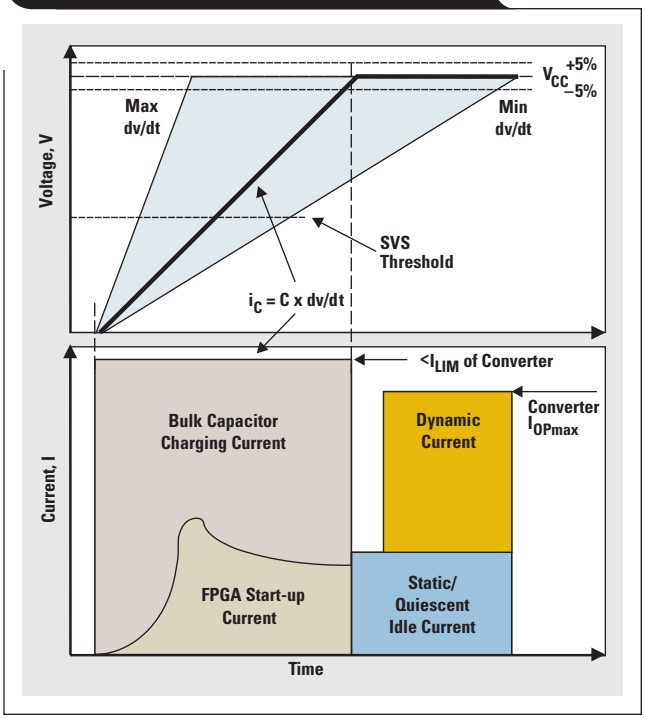

During startup, the capacitors on the power rail are charging due to current flowing into them. In addition to this current, some older, multi-rail FPGA and DSP ICs require large current surges during startup due to bus contention or other undesirable transistor states. Maximum voltage rail ramp time requirements (causing minimum dv/dt slew rates and slower start-up times) place a limit on the length of time that subcircuits within the IC are held below their operating voltage. In some older digital ICs, slower ramping core voltage rails reduced the start-up surge current into the IC. As illustrated in the bottom graph of Figure 5, many of the newer FPGAs and DSPs have resolved such start-up issues (e.g., by sequentially starting up different sections of the IC), and their start-up currents are well below the expected dynamic current levels. So for the most recent FPGAs and DSPs, the current for charging bulk capacitance in the rail's decoupling network dominates the start-up inrush current. Using $i_C = C \times dv/dt$, where C is the bulk capacitance and dv/dt is the selected ramp rate, we can easily see that the peak current could potentially be larger than the digital IC's dynamic current or the maximum operating current of that rail's converter. Minimum voltage ramp time requirements (resulting in maximum dv/dt slew rates and fast start-up times) for FPGA or DSP voltage rails arise for a variety of reasons, including IC testing limitations. Regardless of whether or not minimum ramp time requirements are present, the POL converter should be soft started to ensure that the inrush current for charging bulk capacitance does not exceed the current limit of the converter, potentially causing it to start up improperly as it protects itself from damage from the overcurrent. Part 2 of this series shows examples of two common start-up problems seen when converters/regulators are used and explains how to use soft start to control the voltage rail ramp rate and fix these problems.

Methodology

Compiling the previous recommendations, Figure 6 provides a simple method for powering up each of the digital ICs in the simplified application in Figure 1.

The choice of linear or switching regulators depends on the rail's power consumption (i.e., the efficiency needed for power dissipation) and/or the application needs (i.e., low-noise requirements). This method assumes that the application is configured to allow the FPGA or DSP to be powered up before the other (i.e., sequentially). The power solution in Figure 6 uses sequential sequencing of each digital device by tying the output voltage, if high enough, or the integrated or external SVS of the first rail

Figure 6. Block diagram of power solution

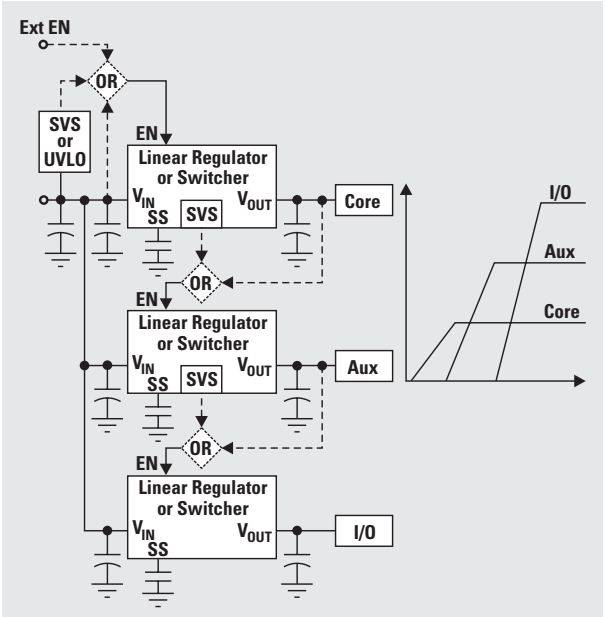

to the enable pin of the next rail. Although the example in Figure 6 shows the default sequencing order of core voltage, auxiliary voltage(s), and I/O voltage, the exact sequencing order should follow the FPGA or DSP manufacturer's recommendations, if any. In some cases, the recommended sequencing order may not be required for successful power up or to prevent damage, but it may reduce start-up currents. Compared to the other sequencing methods, sequential sequencing is easier to implement and further minimizes start-up currents by staggering power-rail startup. Using a POL converter with a soft start that allows the power rail's rise time to be controlled is also highly recommended and may even be required if the FPGA or DSP has minimum rise- and fall-time requirements for one or more power rails. The external SVS with controlled delay is necessary for many DSPs requiring an external RESET signal. The last, but arguably the most important, step in the power-supply design process is designing for load transients.

Conclusion

Designing a multi-rail power solution involves more than simply minimizing size and cost. System level and/or application requirements, such as noise sensitivity or ambient temperature, may require or prevent the use of a low-cost linear regulator. Considering how the rails interact during startup is also critical for a robust system. How to choose the right converter and then design it for a multi-rail application is covered in Part 2 of this series.

Related Web site

power.ti.com

TPS79918 RF LDO supports migration to StrataFlash® Embedded Memory (P30)

By Michael Day (Email: m-day@ti.com)
Power Management Products/Portable Power dc/dc Applications

Introduction

The Texas Instruments (TI) TPS79918 low dropout (LDO) linear regulator provides the performance required to power the new Intel® StrataFlash® Embedded Memory (P30). Intel is migrating from its third-generation 180-nm StrataFlash Embedded Memory (J3) to its fourth-generation 130-nm StrataFlash Embedded Memory (P30). This migration from J3 to P30 memory allows systems to operate at a lower overall current consumption because the P30 V_{CC} voltage requirement has dropped to 1.8 V. Intel Application Note 812[1] recommends using an LDO linear regulator to provide the new 1.8-V voltage rail.

TPS79918 features

The TPS79918 LDO offers exceptional electrical performance in a small package. It provides a high (greater than 66-dB) power supply rejection ratio (PSRR), low noise, fast startup, and excellent line and load transient response. It consumes only 40 µA of ground current at full load. No input capacitor is required for operation. The TPS79918 is stable with a small ceramic output capacitor and uses an advanced BiCMOS fabrication process to yield a 110-mV dropout voltage at the full 200-mA output load. Its precision voltage reference and feedback loop achieve a 2% overall accuracy over all load, line, process, and temperature variations. The addition of an optional noise-reduction capacitor provides an extremely low-noise output voltage. It is fully specified from T_J = –40°C to +125°C and is offered in three packages: ThinSOT-23, wafer chip-scale (WCSP), and 2 × 2 SON-6. The difference between the three packages is size and power dissipation. If we assume a 60°C ambient temperature, the power dissipation for the three packages is 142 mW, 156 mW, and 526 mW, respectively. That corresponds to a maximum continuous load current of 119 mA, 130 mA, and 438 mA, respectively.

Figure 1 shows the TPS79918 internal block diagram. The active-high enable signal turns the device on after the input voltage goes above the undervoltage lockout (UVLO) level. UVLO ensures that the input voltage is available, which provides a clean start-up waveform. The IC features current-limit and thermal-shutdown protection in case of a system failure. The internal 500-kΩ resistor, along with an external noise-reduction capacitor on the NR pin, provides an RC filter to reduce the bandgap noise, resulting in a low-noise output voltage. The quickstart switch initially shorts out this resistor to provide a fast 40-µs start-up time. This IC has excellent transient response enhanced by the Overshoot Detect circuit that during a heavy load actively pulls the output voltage down to a light-load transient.

Figure 1. TPS79918 internal block diagram

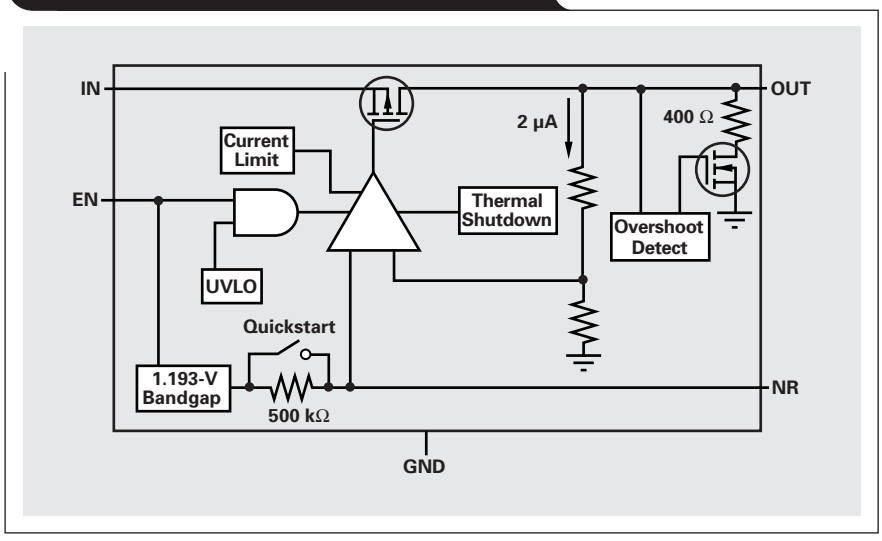

Figure 2. TPS79918 load transient

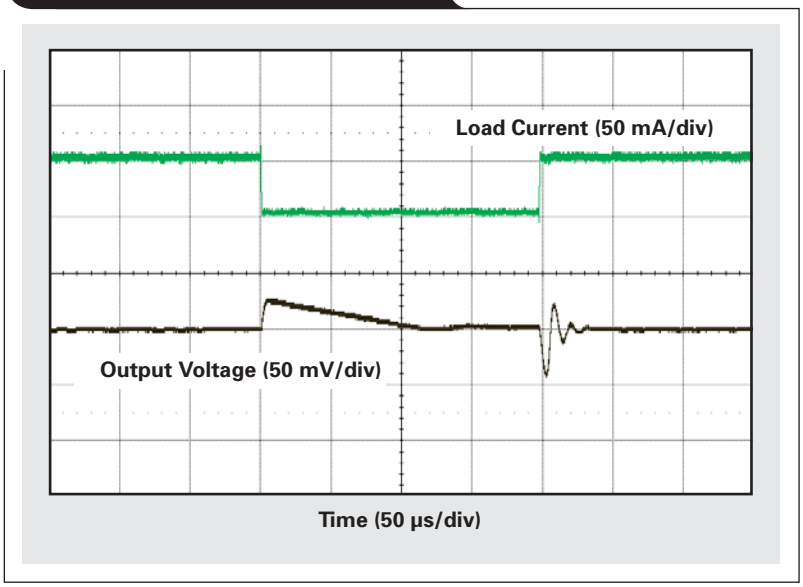

Intel StrataFlash Embedded Memory (P30) requirements

Migrating existing J3 memory systems to the new P30 memory requires regulating the existing 3.0-V core voltage down to 1.8 V. The 1.8-V core voltage requires regulation within ±0.1 V, which is met by the TPS79918's tight 2% voltage tolerance. The LDO performs well with the load transients caused by rising and falling edges of CE# and OE#. Figure 2 shows that even with a 1- to 50-mA load transient, the output voltage drops only 40 mV. This performance meets the P30 memory input voltage requirement. Taking system-level requirements into consideration, Intel's application note cautions the designer to consider both standby current and active current. Figure 3 shows that, unlike some LDOs, the TPS79918's low quiescent current is constant over changes in input voltage and load current. The low quiescent current extends battery life in portable applications.

Figure 3. TPS79918 quiescent current vs. line and load

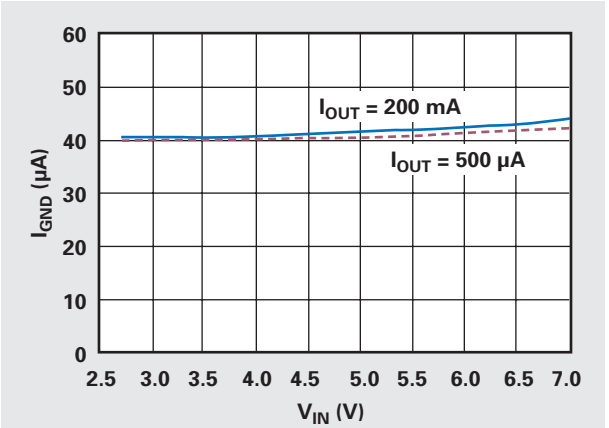

Figure 4 shows the simplicity of a TPS79918 powering a P30 memory module. The only required components are the LDO itself and its output capacitor. The input capacitor is required only if the 3.0-V output capacitor is not physically located near the LDO. The noise-reduction capacitor is needed only to power noise-sensitive applications, so it is not required in this application.

Intel specifically recommends the TI LDOs in Table 1 for its StrataFlash Embedded Memory (P30). These LDOs meet the electrical and temperature requirements for the new P30 memory.

Table 1. Recommended TI LDOs for Intel StrataFlash Embedded Memory (P30)

DEVICE	V_O (V)	I_O (mA)	I_Q (µA)	PSRR (dB)	PACKAGE	DIMENSIONS (mm)
TPS79918DDC	1.8	200	40	70	TSOT-23	2.9 x 2.8 x 1.0
TPS79918DRV	1.8	200	40	70	2 x 2 SON-6	2.0 x 2.0 x 0.75
TPS79918YZU	1.8	200	40	70	WCSP	1.36 x 1.00 x 0.62

Conclusion

The TI TPS79918 LDO is ideally suited to power Intel's next-generation StrataFlash Embedded Memory (P30). The three packaging options allow the designer to optimize the solution size while the LDO's electrical performance meets the system-level requirements for low standby current, transient response, and peak current.

Reference

1. "Migration Guide for Intel StrataFlash® Memory (J3) to Intel StrataFlash® Embedded Memory (P30)," Intel Application Note 812, June 3, 2005, www.intel.com/design/flcomp/applnots/306667.htm

Related Web sites

power.ti.com
www.intel.com

Figure 4. Typical StrataFlash configuration with LDO

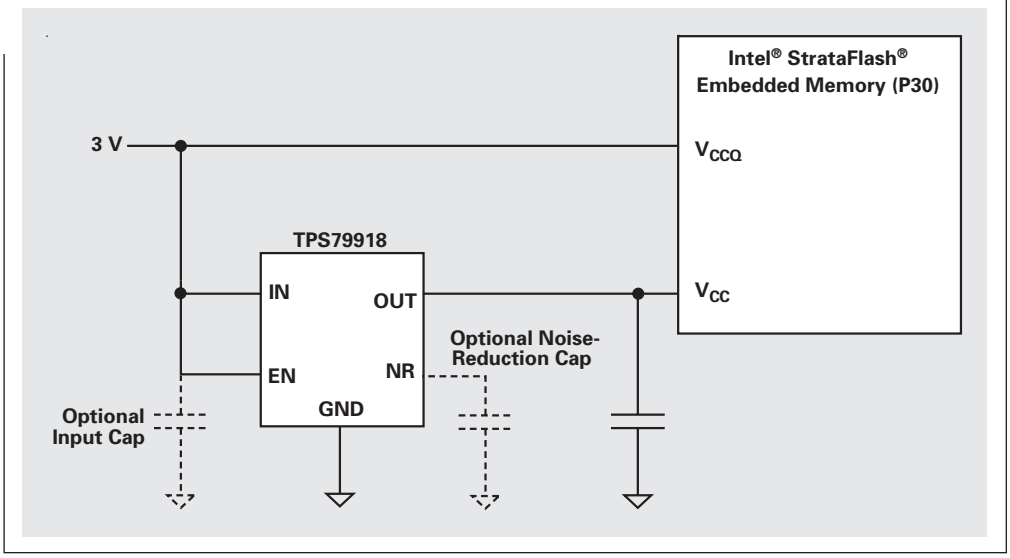

Practical considerations when designing a power supply with the TPS6211x

By Jeff Falin, *Power Management Products/Portable Power dc/dc Applications* (Email: j-falin1@ti.com),
and Bill Johns, *Power Management Products/Portable Power dc/dc Applications* (Email: w-johns2@ti.com)

The TPS6211x dc/dc converter is a synchronous buck converter capable of input voltages up to 17 V, output voltages from 1.2 to 16 V, and output currents up to 1.5 A. The device efficiently steps down 2-cell Li-ion or lead-acid batteries, or 12- to 15-V system rails, to 5 V, 3.3 V, or lower. A typical application circuit is shown in Figure 1.

This article provides an example application and summarizes the key features of designing with the TPS6211x buck converter.

Figure 1. Typical application

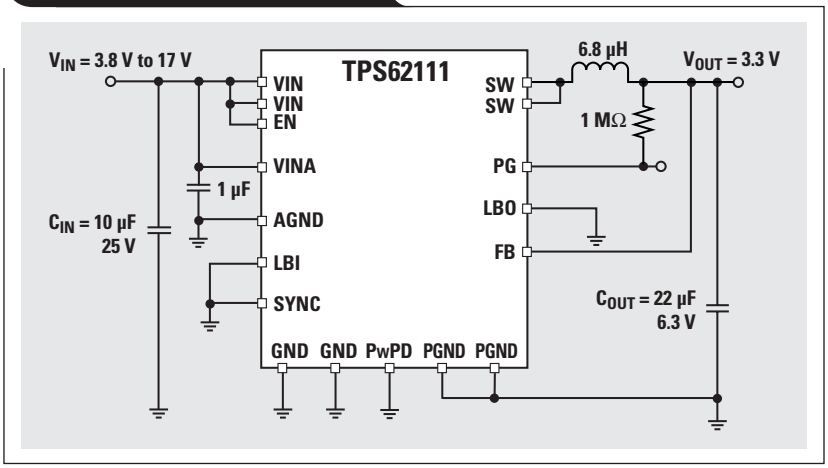

External component sizing

Integrated FETs and internal compensation reduce the required number of external parts to very few. Assuming that the integrated output-voltage and input-voltage supervisory circuits are not used, the fixed-output-voltage (5-V or 3.3-V) versions need only four components: an input capacitor for the power switches, an input filter capacitor for the analog input pin, an output capacitor (C), and a power inductor (L). The device has been internally compensated to work with an L of 6.8 µH and a C of at least 22 µF. Other values of L and C will provide satisfactory operation, and increasing either or both values will reduce output ripple. It is recommended that the LC product remain close to the minimum of 6.8 µH × 22 µF for stability. Two additional feedback resistors and a feed-forward capacitor are required for the adjustable version, which provides output voltages from 1.2 to 16 V. The output voltage supervisory circuit (called power good, or PG, output) provides a logic high when the output is above approximately 98.6% of its nominal voltage. The low-battery indicator is a stand-alone supervisory circuit, active only after EN is pulled high, with a low-battery input (LBI) that is intended to monitor battery input voltage and a low-battery output (LBO) that provides a logic low when the input voltage drops below a certain voltage. LBI requires two external feedback resistors to set the trip point and, if not used, should be pulled low as shown in Figure 1. Both PG and LBO are open-drain outputs providing for maximum user flexibility and therefore require pull-up resistors if used but can be left floating if not used.

Output voltage regulation versus output current

The device's proprietary control architecture ensures that the output voltage maintains ±3% output voltage regulation (excluding feedback resistor tolerances) across input line and output load changes. For 17 V > V_{IN} > 6 V, the IC's high-voltage PMOS FET has plenty of gate drive to provide ±3% regulation from no load to 1.5 A. At lower input voltages, the IC still provides ±3% regulation from no load up to 1.2 A for 6 V > V_{IN} > 4.3 V; 500 mA for 4.3 V > V_{IN} > 3.5 V; and 300 mA for 3.5 V > V_{IN} > 3.1 V. Using a high-side PMOS FET allows 100% duty-cycle operation, during which time the output is tied directly to a nearly depleted battery to maximize battery life.

Assuming that the recommended output-filter product is used, the device is compensated to provide excellent load-transient performance. Therefore, fast changes in output load do not result in significant output-voltage droop before the control loop responds. As shown in Figure 2, the output voltage droops less than 100 mV during a 1-A

Figure 2. The 1-A load transient going from 12 to 3.3 V when L = 6.8 µH and C = 22 µF

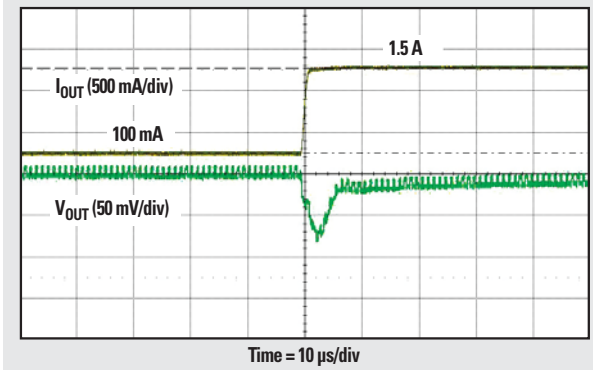

load transient (in forced pulse-width modulation [PWM] mode) with a small 22-µF output capacitor. A larger output capacitor reduces the droop even further.

PWM versus power-save mode

Efficiency varies inversely with the input to output voltage differential for all converters. In addition, efficiency drops dramatically at light loads for fixed-frequency PWM converters due to the IC's quiescent current and switching losses dominating the output power being provided. To improve efficiency at light load, the TPS6211x has a pulse-frequency modulation (PFM), or power-save, mode that is enabled by tying the SYNC pin low. Efficiency curves for V_{OUT} = 3.3 V and for various input voltages with power-save mode enabled are shown in Figure 3.

Figure 3. Efficiency of 3.3-V fixed-output version of TPS62111 with power-save mode enabled

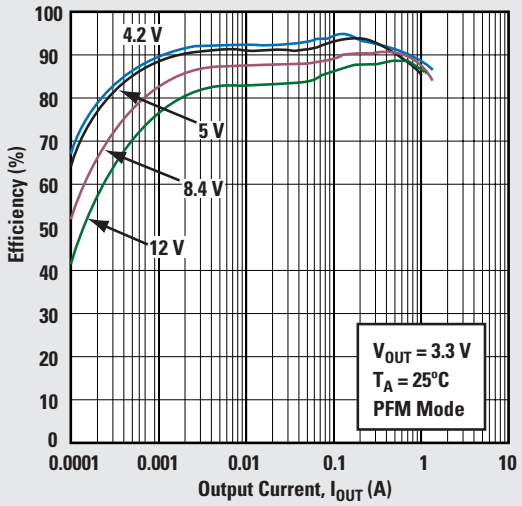

In power-save mode, the TPS6211x's quiescent current is reduced, and the switches turn on only long enough to keep the output voltage in regulation for a given load current. This type of operation causes the output-voltage ripple to be slightly higher (up to 1% of the output voltage) and to vary in frequency until being transitioned back to fixed-frequency PWM mode at currents above approximately 280 mA for V_{IN} > 7 V. In power-save mode there is no minimum duty cycle. Low $r_{DS(ON)}$ FETs with copper overlay reduce I^2R losses and keep efficiency high at heavy load as well.

The TPS6211x can be forced into fixed-frequency PWM mode, resulting in extremely low-amplitude, fixed-frequency output ripple (typically < 10 mV_{PP}) across the entire load range, either by pulling SYNC high to V_{IN} or by tying it to an external oscillator. Pulling SYNC high results in the 1-MHz, fixed-frequency operation. The device can be synchronized through the SYNC pin to an external oscillator with a frequency between 800 kHz and 1.4 MHz. In addition, since the synchronous FET allows current to flow back to the input, there is always current flowing though the inductor, keeping it in continuous conduction mode over the entire load range. This eliminates the ringing that occurs at the switch node after other converters enter discontinuous mode. Therefore, at the expense of slightly lower efficiency at light load, forcing the TPS6211x into PWM mode provides extremely low-noise operation.

Output power

The TPS6211x's 4-mm × 4-mm QFN package with PowerPAD™ has a junction-to-pad thermal resistance ($R_{\Theta JP}$) of 8°C/W and a junction-to-ambient thermal resistance ($R_{\Theta JA}$) of approximately 40°C/W for a high-K board with no air flow. With power dissipation computed as $P_{Dmax} = (T_{Jmax}–T_A)/R_{\Theta JA}$, maximum output power is computed as $P_{OUTmax} = P_{Dmax}/(1/\eta – 1)$, where T_{Jmax} = 125°C and η is the expected efficiency. For example, with typical efficiencies of 85% and 89% at I_{OUT} = 1.5 A, the 3.3- and 5-V fixed-output versions can easily provide 4.95 and 7.5 W, respectively, from a 12-V input rail at 85°C ambient temperature.

Typical applications

The TPS6211x family of buck converters is ideally suited for a wide range of applications. For example, in the computing world, the converter's fast transient response is useful for stepping down 12-V input rails to the 5- or 3.3-V (or lower) rails that power fast-switching digital ICs in a server/workstation or in a personal printer. Consumer applications such as set-top boxes, which typically have 12-V ac/dc power supplies, and portable DVD players, with either 9-V/12-V wall adapters or dual Li-ion (8.4-V nominal) input power, have traditionally needed their input rails stepped down to a 5- and/or 3.3-V midrange bus voltage for I/O power. Further down-conversion of these bus voltages by additional dc/dc switchers and/or linear regulators to power the lower-voltage ICs in the box is also required. However, with its high-efficiency, low-noise operation and low minimum duty cycle, the TPS6211x can power the lower-voltage (e.g., 2.5-V, 1.x-V) processors and peripherals directly, thereby eliminating the midrange bus and improving efficiency and battery life, if applicable. Since 12- and 15-V rails are common in industrial applications, this converter, with its low-noise mode and synchronization feature, could be useful as a point-of-load regulator powering a data acquisition system. Lastly, 2-cell Li-ion and lead-acid battery-powered systems such as battery-backup or alarm systems use the LBI/LBO feature, high efficiency, and 100% duty cycle mode to provide maximum operation time by extending battery life.

Related Web sites

power.ti.com
www.ti.com/sc/device/TPS62111

High-speed notch filters

By Bruce Carter (Email: r-carter5@ti.com)
Low-Power Wireless Applications

Introduction

Active notch filters have been used in the past for applications like elimination of 50- and 60-Hz hum components. They have proven to be somewhat problematic from the standpoints of center frequency (f_0) tuning, stability, and repeatability. The advent of high-speed amplifiers opens the possibility of higher-speed notch filters—but are they actually producible? This article will show what is presently possible and what design trade-offs a designer will face with real-world components.

As a review, the reader should remember some characteristics of the notch filter:

- The depth of the notch obtainable in simulations like that shown in Figure 1 is *not* the depth that can be achieved with real-world components. The best that the designer can hope for is 40 to 50 dB.
- Instead of focusing on notch depth, the designer should focus on center frequency and Q. The Q for a given notch filter is the –3-dB point, *not* the notch depth or a point 3 dB above the notch depth, as shown in Figure 2.

Remember that the designer's objective is not a notch filter but the rejection of a specific interfering frequency. Any filter that does not reject that interfering frequency because it misses the frequency or has too little rejection at that frequency is not much use.

The best way to avoid missing the interfering frequency is to select the best values of R and C from the start. The RC Calculator under "Filter Design Utilities" in Reference 1 should be used to find the correct values of R_0 and C_0 for the circuits in the following discussion.

Topology

A number of notch-filter topologies were explored. Some design goals are a topology that:

- produces a notch (as opposed to band rejection);
- uses a single op amp;
- can be easily tuned with independent adjustments for center frequency and Q;
- can operate from a single-supply voltage; and
- can be adapted to fully differential op amps.

Unfortunately, it was not possible to achieve all of these, although some desirable circuits can be constructed that can meet some of these goals.

Figure 1. Simulated notch depth

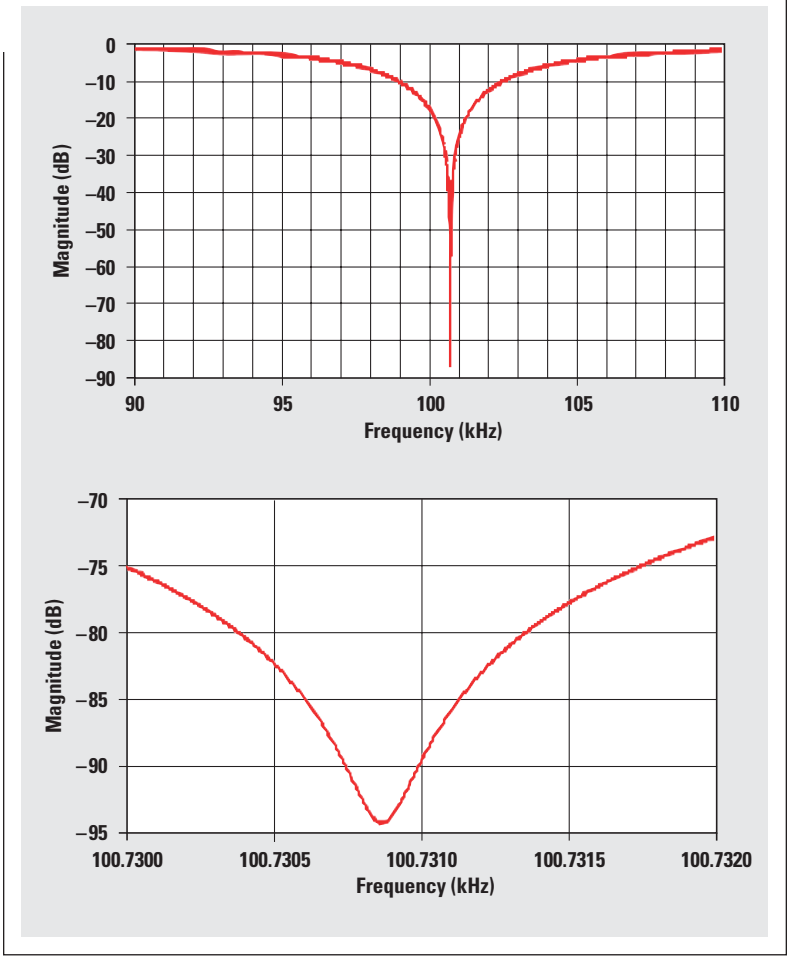

Figure 2. The Q of a notch filter

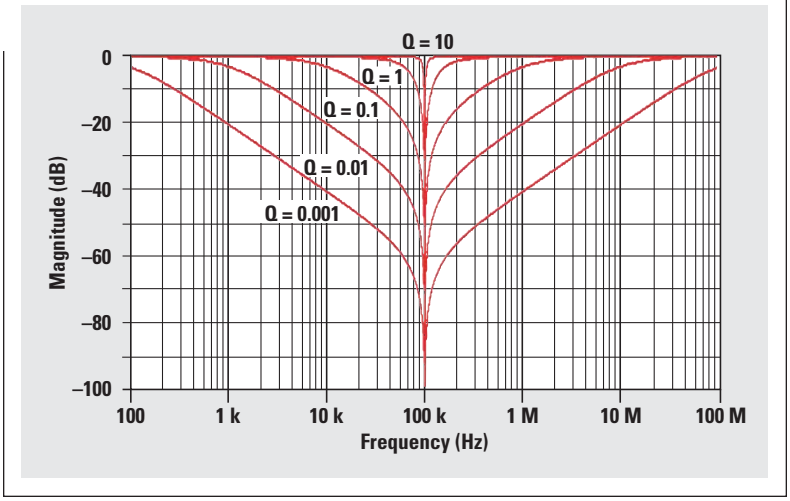

Twin-T notch filter

The twin-T topology of Figure 3 deserves an honorable mention here, because a notch filter can be implemented with a single op amp. It is not as flexible as one would hope, because the center frequency is not easily adjustable. Trimming the center frequency involves simultaneous adjustment of the three R_0 resistors. This is a concern because triple potentiometers are large, expensive, and may not track very well—especially the section that has to be one-half the value of the other two. Mismatches in the R_0 resistors will very quickly erode notch depth to less than 10 dB.

The circuit has some other disadvantages as well:

- It requires six high-precision components for tuning, and two of those are ratios of the others. If the designer wants to get away from ratios, eight precision components are required. $R_0/2$ = two R_0 in parallel, and $2 \times C_0$ = two C_0 in parallel.
- The twin-T topology is not easily adaptable to single-supply operation and cannot be used with a fully differential amplifier.
- The spread of resistor values becomes large due to the requirement of $R_Q << R_0$. The spread of the resistor values has a bearing on the depth of the notch and on center frequency.

Nevertheless, for applications where only a single op amp can be used, the twin-T topology is quite usable if the designer matches components or buys very high-precision components.

Fliege notch filter

The Fliege notch topology is shown in Figure 4. The advantages of this circuit over the twin-T are as follows:

- Only four precision components—two Rs and two Cs—are required for tuning the center frequency. One nice feature of this circuit is that slight mismatches of components are okay—the center frequency will be affected, but not the notch depth.
- The Q of the filter can be adjusted independently from the center frequency by using two noncritical resistors of the same value.

Figure 3. Twin-T notch filter

$$f_0 = \frac{1}{2\pi R_0 C}$$

$$R_Q << R_0$$

$$Q = \frac{R_Q 2}{4 \times R_Q 1}$$

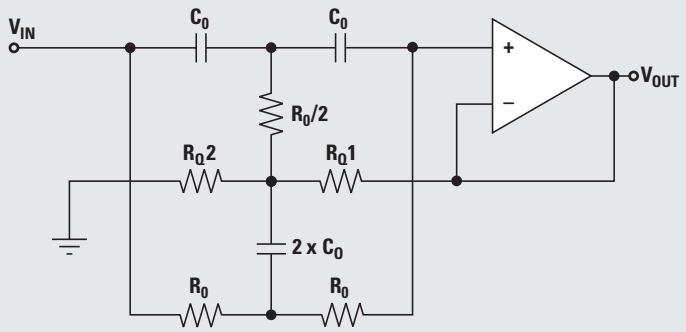

Figure 4. Fliege notch filter

$$f_0 = \frac{1}{2\pi R_0 C_0}$$

$$Q = \frac{R_Q}{2 \times R_0}$$

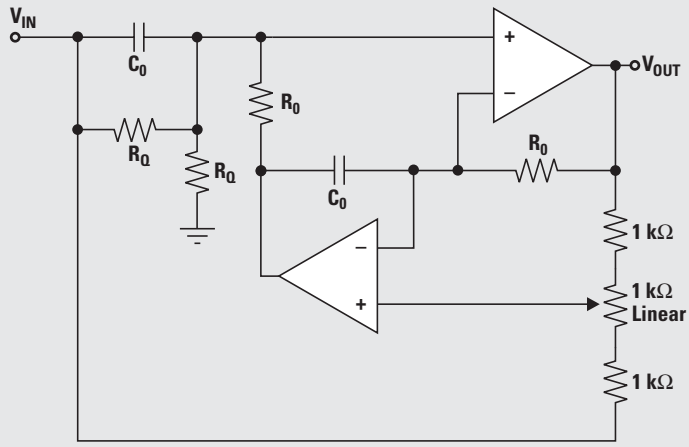

Table 1. Component values for the Fliege notch filter

Q	1 MHz			100 kHz			10 kHz		
	R_0 (kΩ)	C_0 (pF)	R_Q (kΩ)	R_0 (kΩ)	C_0 (nF)	R_Q (kΩ)	R_0 (kΩ)	C_0 (nF)	R_Q (kΩ)
100	1.58	100	316	1.58	1	316	1.58	10	316
10	1.58	100	31.6	1.58	1	31.6	15.8	1	316
1	1.58	100	3.16	1.58	1	3.16	15.8	1	31.6

- The center frequency of the filter can be adjusted over a narrow range without seriously eroding the depth of the notch.

Unfortunately, this circuit uses two op amps instead of one, and it cannot be implemented with a fully differential amplifier.

Simulations

Simulations were first performed with ideal op amp models. Real op amp models were later used, which produced results similar to those observed in the lab. Table 1 shows the component values that were used for the schematic in Figure 4. There was no point in performing simulations at or above 10 MHz because lab tests were actually done first, and 1 MHz was the top frequency at which a notch filter worked.

A word about capacitors: Although the capacitance is just a value for simulations, actual capacitors are constructed of different dielectric materials. For 10 kHz, resistor value spread constrained the capacitor to a value of 10 nF. While this worked perfectly well in simulation, it forced a change from an NPO dielectric to an X7R dielectric in the lab—with the result that the notch filter completely lost its characteristic. Measurements of the 10-nF capacitors used were close in value, so the loss of notch response was most likely due to poor dielectric. The circuit had to revert to the values for a Q of 10, and a 3-MΩ R_Q was used. For real-world circuits, it is best to stay with NPO capacitors.

The component values in Table 1 were used both in simulations and in lab testing. Initially, the simulations were done without the 1-kΩ potentiometer (the two 1-kΩ fixed resistors were connected directly together and to the noninverting input of the bottom op amp). Simulation results are shown in Figure 5.

There are actually nine sets of results in Figure 5, but the curves for each Q value overlie those at the other frequencies. The center frequency in each case is slightly above a design goal of 10 kHz, 100 kHz, or 1 MHz. This is as close as a designer can get with a standard E96 resistor and E12 capacitor. Consider the case of 100 kHz:

$$f_0 = \frac{1}{2\pi R_0 C_0} = \frac{1}{2\pi \times 1.58\ \text{k}\Omega \times 1\ \text{nF}} = 100.731\ \text{kHz}$$

A closer combination exists if E24 sequence capacitors are available:

$$f_0 = \frac{1}{2\pi R_0 C_0} = \frac{1}{2\pi \times 4.42\ \text{k}\Omega \times 360\ \text{pF}} = 100.022\ \text{kHz}$$

The inclusion of E24 sequence capacitors can lead to more accurate center frequencies in many cases, but procuring the E24 sequence values is considered an expensive (and

Figure 5. Simulation results before tuning

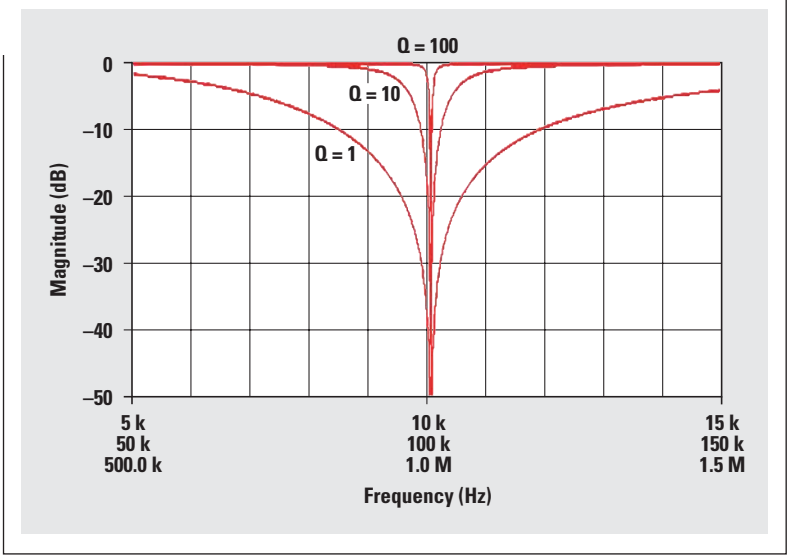

unwarranted) expenditure in many labs. While it may be easy to specify E24 capacitor values in theory, in practice many of them are seldom used and have long lead times associated with them.

There are easier alternatives to selecting E24 capacitor values. Close examination of Figure 5 shows that the notch misses the center frequency by only a small amount. At lower Q values, there is still substantial rejection of the desired frequency. If the rejection is not sufficient, then it becomes necessary to tune the notch filter.

Again considering the case of 100 kHz, we see that the response near 100 kHz is spread out in Figure 6. The family of curves to the left and right of the center frequency (100.731 kHz) represents filter response when the 1-kΩ potentiometer is inserted and adjusted in 1% increments. When the potentiometer is exactly in the middle, the notch filter rejects frequencies at the exact center frequency. The depth of the simulated notch is actually on the order of 95 dB, but that is not going to happen in the real world. A 1% adjustment of the potentiometer puts a notch that is greater than 40 dB right on the desired frequency. Again, this is best-case with ideal components, but lab results are close at low frequencies (10 and 100 kHz).

Figure 6 shows that it is important to get close to the correct frequency with R_0 and C_0 from the start. While the potentiometer can correct for frequency over a broad range, the depth of the notch degrades. Over a small range (±1%), it is possible to get a 100:1 rejection of the undesirable frequency; but over a larger range (±10%), only a 10:1 rejection is possible.

Figure 6. Tuning for center frequency

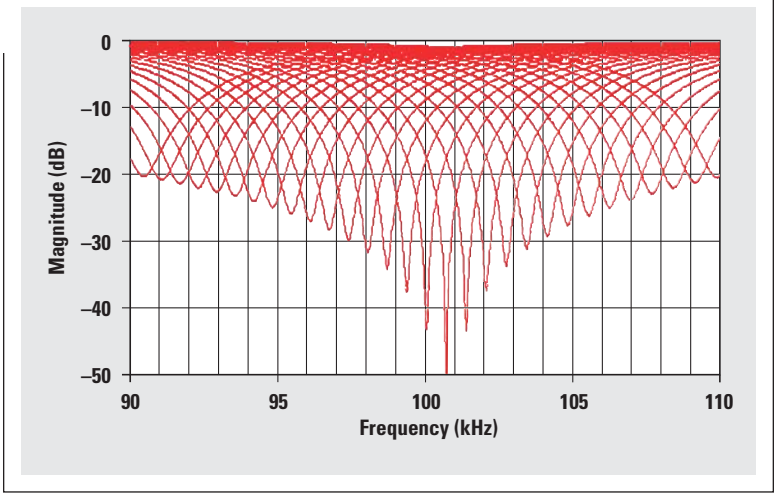

Lab results

A THS4032 evaluation board was used to construct the circuit in Figure 4. Its general-purpose layout required only three jumpers and one trace cut to complete the circuit. The component values in Table 1 were used, starting with those that would produce 1 MHz. The intention was to look for bandwidth/slew-rate restrictions at 1 MHz and test at lower or higher frequencies as necessary.

Results at 1 MHz

Figure 7 shows that there are some very definite bandwidth and/or slew-rate effects at 1 MHz. The response curve at a Q of 100 shows barely a ripple where the notch should be. At a Q of 10, there is only a 10-dB notch, and a 30-dB notch at a Q of 1. Apparently notch filters cannot achieve as high a frequency as one would hope, but the THS4032 is only a 100-MHz device. It is reasonable to expect better performance from parts with a greater unity-gain bandwidth. Unity-gain stability is important, because the Fliege topology has fixed unity gain.

If the designer wishes to estimate what bandwidth is required for a notch at a given frequency, a good place to

Figure 7. Lab results at 1 MHz

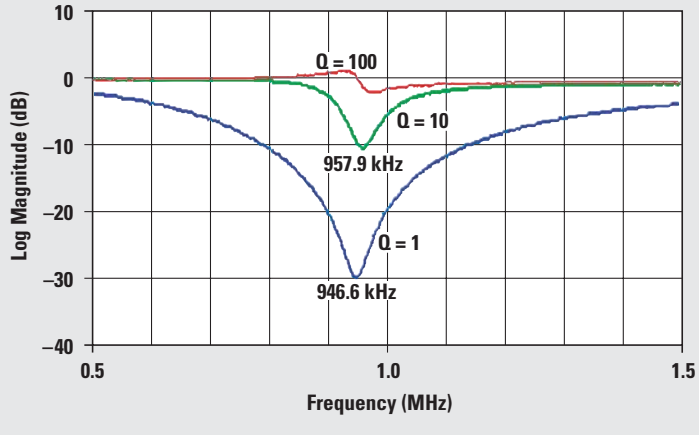

start is the gain/bandwidth product given in the datasheet, which should be 100 times the center frequency of the notch. Additional bandwidth will be required for higher Q values. There is a slight frequency shift of the notch center as Q is changed. This is similar to the frequency shift seen for bandpass filters. The frequency shift is less for notch filters centered at 100 kHz and 10 kHz, as shown in Figure 8 and later in Figure 10.

Results at 100 kHz

Component values from Table 1 were then used to create 100-kHz notch filters with different Qs. The results are shown in Figure 8. It is immediately obvious that viable notch filters can be constructed with a center frequency of 100 kHz, although the notch depth appears to be less at higher values of Q.

Remember, though, that the design goal here is a 100-kHz—not a 97-kHz—notch. The component values selected were the same as for the simulation, so the notch center frequency should theoretically be at 100.731 kHz; but the difference is explained by the parts used in the lab. The mean value of the 1000-pF capacitor stock was 1030 pF, and of the 1.58-kΩ resistor stock was 1.583 kΩ. When the center frequency is calculated with these values, it comes out to 97.14 kHz. The actual components, however, could not be measured (the board was too fragile).

As long as the capacitors are matched, it would be possible to go up a couple of standard E96 resistor values to get closer to 100 kHz. Of course, this is probably not an option in high-volume manufacturing, where 10% capacitors could come from any batch and potentially from different manufacturers. The range of center frequencies will be determined by the tolerances of R_0 and C_0, which is not good news if a high Q notch is required. There are three ways of handling this:

- Purchase higher-precision resistors and capacitors;
- lower the Q requirement and live with less rejection of the unwanted frequency; or
- tune the circuit (which was explored next).

At this point, the circuit was modified to have a Q of 10, and a 1-kΩ potentiometer was added for tuning the center frequency (as shown in Figure 4). In real-world design, the potentiometer value selected should slightly more than cover the range of center frequencies possible with worst-case R_0 and C_0 tolerances. That was not done here, as this was an exercise in determining possibilities, and 1 kΩ was the lowest potentiometer value available in the lab.

Figure 8. Lab results at 100 kHz

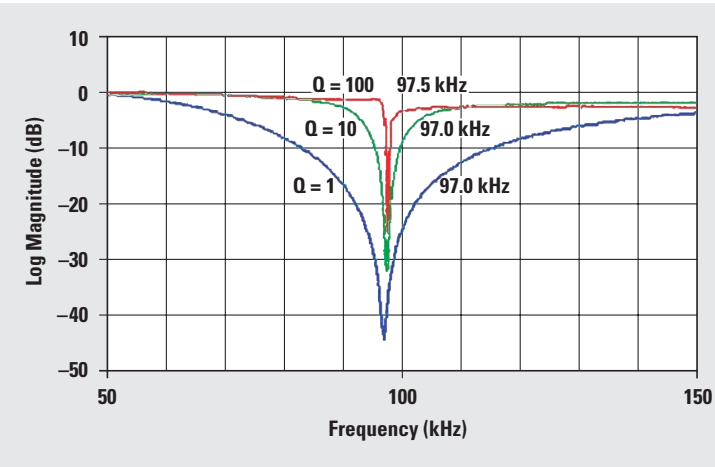

Figure 9. Tuning for exact center frequency

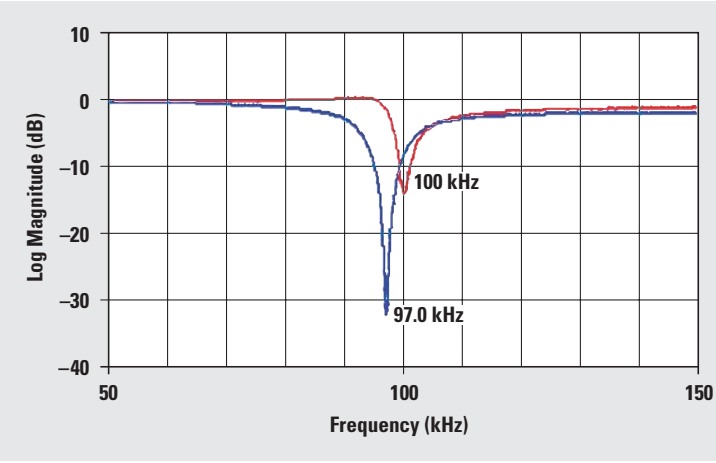

When the circuit was tuned for a center frequency of 100 kHz as shown in Figure 9, the notch depth degraded from 32 dB to 14 dB. Remember that this notch depth could be greatly improved by making the initial f_0 closer to ideal. The potentiometer is meant to tune over only a small range of center frequencies. Still, a 5:1 rejection of an unwanted frequency is respectable and may be sufficient for some applications. More critical applications will obviously need higher-precision components.

Op amp bandwidth limitations, which will also degrade the tuned notch depth, may also be keeping the notch depth from being as low as possible. With this in mind, the circuit was retuned for a center frequency of 10 kHz.

Results at 10 kHz

Figure 10 shows that the notch depth for a Q of 10 has increased to 32 dB, which is about what one would expect from a center frequency 4% off from the simulation (Figure 6). The op amp was indeed limiting the notch depth at a center frequency of 100 kHz! A 32-dB notch is a rejection of 40:1, which is quite good.

So even with components that produced an initial 4% error, it was possible to produce a 32-dB notch at the desired center frequency. The bad news is that to escape op amp bandwidth limitations, the highest notch frequency possible with a 100-MHz op amp is somewhere between 10 and 100 kHz. In the case of notch filters, "high-speed" is therefore defined as being somewhere in the tens or hundreds of kilohertz.

A good application for 10-kHz notch filters is AM (medium-wave) receivers, where the carrier from adjacent stations produces a loud 10-kHz whine in the audio, particularly at night. This can really grate on one's nerves when listening is prolonged. Figure 11 shows the received audio spectrum of a station before and after the 10-kHz notch was applied. Note that the 10-kHz whine is the loudest portion of the received audio (Figure 11a), although the human ear is less sensitive to it. This audio spectrum was taken at night on a local station that had two strong stations on either side. FCC regulations allow for some variation of the station carriers. Therefore, slight errors in carrier frequency of the two adjacent stations will make the 10-kHz tones heterodyne, increasing the unpleasant listening sensation. When the notch filter is applied (Figure 11b), the 10-kHz tone is reduced to the same level as that of the surrounding modulation. Also visible on the audio spectrum are 20-kHz carriers from stations two channels away and a 16-kHz tone from a transatlantic station. These are not a problem, because they are attenuated substantially by the receiver IF. A frequency of 20 kHz is inaudible to the vast majority of people in any event.

Figure 10. Lab results at 10 kHz*

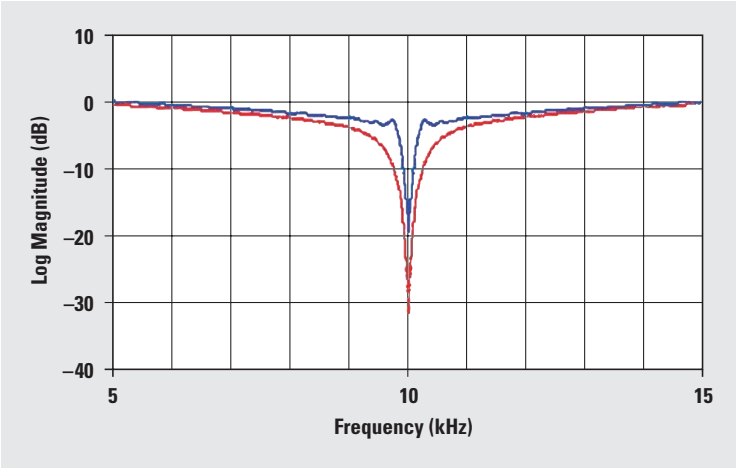

*Some artistic liberties were taken with this plot. The laboratory instrument displays values only down to 10 kHz, so the left-hand portion of the plot is a mirror image of the right-hand portion. The laboratory instrument also has some roll-off at frequencies below 100 kHz, which was artistically eliminated from this plot.

Figure 11. AM reception with and without notch filter

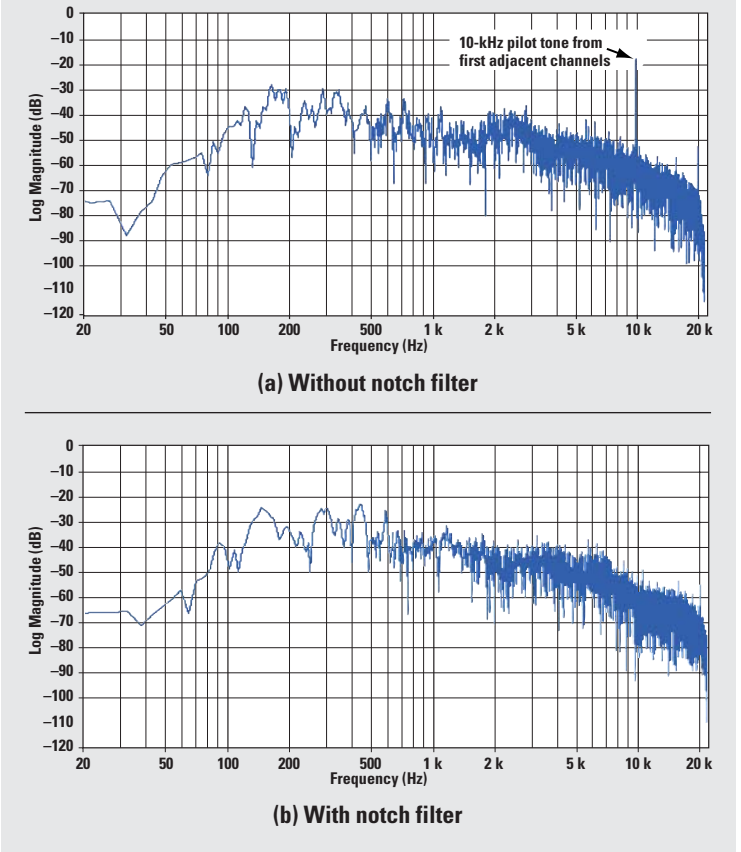

Figure 12 shows the same spectrum on a waterfall diagram. In this case, the sample window is widened, and the 10-kHz carrier interference is shown as a string of peaks that vary in amplitude. When the notch is applied, the 10-kHz peaks are eliminated, and there is only a slight ripple in the received audio where 10 kHz has been notched out.

For European readers who want to have a more pleasing medium-wave listening experience, the component values are C_0 = 330 pF, R_0 = 53.6 kΩ, and R_Q = 1 MΩ. Shortwave listeners will benefit from a two-stage notch filter, one stage being the 10-kHz previously described, and the other stage being a 5-kHz notch filter with component values of C_0 = 270 pF, R_0 = 118 kΩ, and R_Q = 2 MΩ.

Figure 12. Heterodyning and notch filter effects

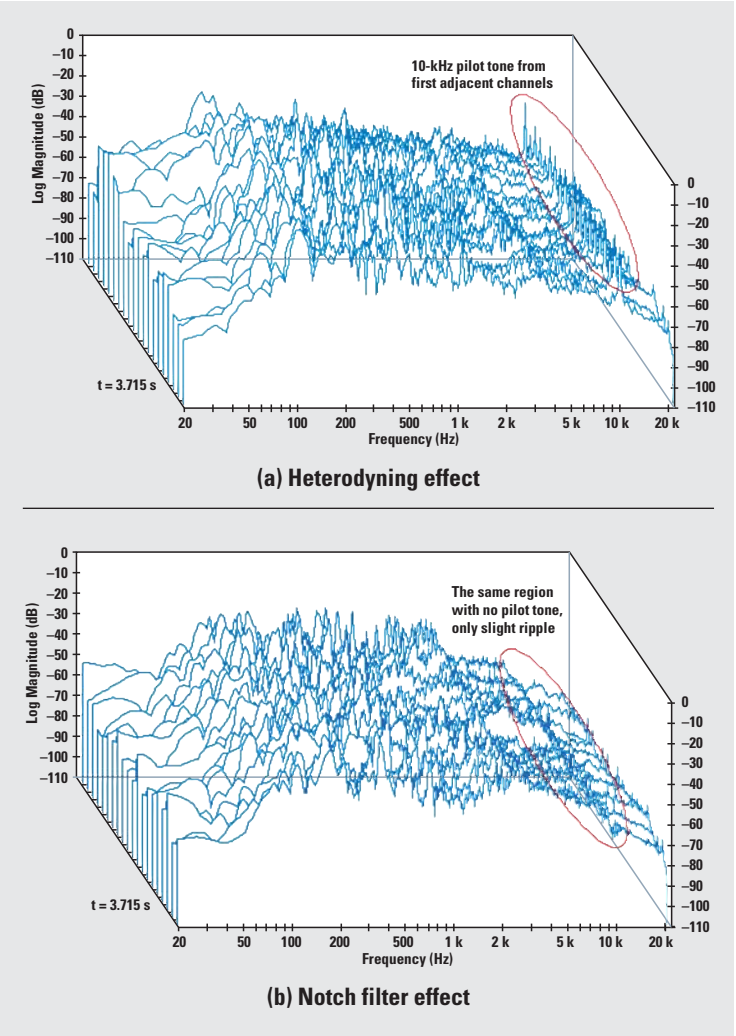

Applicability

Although testing described in this article was performed on the THS4032, the application circuits are usable with all single-ended, unity-gain, voltage-feedback op amps. A key specification is unity-gain bandwidth, which should be from 100 to 1000 times the center frequency. The Fliege notch filter cannot be constructed from current-feedback amplifiers or from fully differential op amps.

Conclusion

High-speed op amps have been used to produce low-pass and high-pass filters up to the tens of megahertz with fairly good success. Narrow band-pass filters and notch filters are much less understood and much more critical applications. While the tolerance of a capacitor might change the cut-off frequency of a low-pass filter or produce ripple in the passband, that same tolerance can produce dramatic changes in the center frequency and notch depth of a notch filter.

With a Fliege notch topology, the number of critical components is reduced to four—two identical Rs and two identical Cs. Fortunately for the designer, there is an inherent matching that occurs when devices are manufactured at the same time, so it is possible to construct notch filters from them even if the tolerance given in the datasheet does not imply matching. There is good, independent control over the center frequency and Q, with the possibility of tuning over a narrow range, which compensates for the initial tolerance errors.

A 1-MHz, Q = 1 notch filter constructed with a 100-MHz op amp showed poor performance at higher values of Q. The same op amp did better at 100 kHz but still showed degradation at higher Q values, particularly when the center frequency was tuned. It was not until the center frequency was decreased to 10 kHz that performance close to simulation results was obtained. Limiting the notch filter to high tens to low hundreds of kilohertz (for faster parts) eliminates many applications. These frequencies, however, represent the state of the art in design for these unusual filters.

Reference

1. "Amplifiers and Linear Engineering Design Utilities," www.ti.com/amplifier_utilities

Related Web sites

amplifier.ti.com
www.ti.com/sc/device/THS4032

Index of Articles

TI Worldwide Technical Support

Internet

TI Semiconductor Product Information Center Home Page
support.ti.com

TI Semiconductor KnowledgeBase Home Page
support.ti.com/sc/knowledgebase

Product Information Centers

Americas

Phone	+1(972) 644-5580
Fax	+1(972) 927-6377
Internet/Email	support.ti.com/sc/pic/americas.htm

Europe, Middle East, and Africa

Phone	
Belgium (English)	+32 (0) 27 45 54 32
Finland (English)	+358 (0) 9 25173948
France	+33 (0) 1 30 70 11 64
Germany	+49 (0) 8161 80 33 11
Israel (English)	180 949 0107
Italy	800 79 11 37
Netherlands (English)	+31 (0) 546 87 95 45
Russia	+7 (0) 95 363 4824
Spain	+34 902 35 40 28
Sweden (English)	+46 (0) 8587 555 22
United Kingdom	+44 (0) 1604 66 33 99
Fax	+(49) (0) 8161 80 2045
Internet	support.ti.com/sc/pic/euro.htm

Japan

Fax	International	+81-3-3344-5317
	Domestic	0120-81-0036
Internet/Email	International	support.ti.com/sc/pic/japan.htm
	Domestic	www.tij.co.jp/pic

Asia

Phone	
International	+886-2-23786800
Domestic	Toll-Free Number
Australia	1-800-999-084
China	800-820-8682
Hong Kong	800-96-5941
India	+91-80-51381665 (Toll)
Indonesia	001-803-8861-1006
Korea	080-551-2804
Malaysia	1-800-80-3973
New Zealand	0800-446-934
Philippines	1-800-765-7404
Singapore	800-886-1028
Taiwan	0800-006800
Thailand	001-800-886-0010
Fax	+886-2-2378-6808
Email	tiasia@ti.com or ti-china@ti.com
Internet	support.ti.com/sc/pic/asia.htm

A120905

SLYT230