



Analog and Interface Guide – Volume 2

A Compilation of Technical Articles and Design Notes



Contents

Techniques To Minimize Noise

Looking Forward To Technology Migration	1
Bypass Capacitors: No Black Magic Here.....	2
The Debounce Debacle	3
ETP-248: Managing Power, Ground And Noise In Microcontroller/Analog Applications	4

Charging Lithium Batteries

Charging High Capacity Lithium Batteries.....	8
Charging Lithium-Ion Batteries: Not All Charging Systems Are Created Equal	10

Operational Amplifiers

Working Your Amplifier Inside The Single-supply Voltage Box.....	14
Every Amplifier Is Waiting To Oscillate And Every Oscillator Is Waiting To Amplify	18
Amplifiers And The SPICE Of Life.....	22

Analog-to-Digital Converters

The Number Of Bits vs LSB Errors	24
Is Your Converter Accurate – Part 2	25
Sometimes Noise Can Be Good	26
Using The Basics For ACD Range Control	27
Turning Nyquist Upside Down By Undersampling.....	28

Miscellaneous

Arming Yourself With nanoWatt Technology Techniques	29
Muxing Your Micro With Style	32
Riding The Crest Factor Wave	33
Designing With Low Dropout Regulators In Embedded Applications.....	34

Techniques To Minimize Noise

Looking Forward To Technology Migration

By Bonnie C. Baker, Microchip Technology Inc.

The industry is laying the groundwork as we speak. The geometries of digital and analog integrated circuit silicon processes are shrinking. As a consequent, IC devices continue to have lower power-supply voltage requirements, silicon areas and prices. I personally have witnessed geometry migration from 1.2 microns all the way down to 0.25 microns. I am finding that a device function manufactured with a 0.7 micron process may not work as well if I change to a 0.25 micron process.

What is that all about?! I theoretically haven't changed anything, but it appears that my circuits present an argument to the contrary. Of course, I have the option of not selecting devices that have migrated to smaller geometries, which is a very bad idea for a forward-thinking engineer. By not participating in this industry trend, I turn down the added benefits of lower power-supply voltages, faster speeds and lower cost. Not very smart!

My only line of defense is to design with the expectation that geometries are going to continue to get smaller. I know that smaller geometry devices are able to support lower power-supply voltages. This helps me with my power consumption issues, but ESD (electrostatic discharge) is less than optimum. Not only are smaller geometry devices less able to absorb high-voltage transients, but they lack robustness around high currents. The manufacture's standards are as high as ever with a 2000 Vrms to 4000 Vrms ESD tolerance (Human Body Model). However, their tests look for catastrophic failures while the end user can experience RAM contamination because of EMI (electromagnetic interference) and EFT (electrical fast transients) signals.

However, there are several areas that I can focus on, such as protection circuits (MOVs, transient suppressors), microcontroller or processor pin protection (I/O, interrupt, reset pins), or firmware recovery techniques (WDT, register refresh), etc. All of these techniques help produce a robust design, but I get the biggest bang-for-the-buck when I optimize my layout.

If you want to join this smaller-geometry migration wave, look for places in your layout where spikes and glitches can enter your, soon to be, sensitive circuits. A great place is to look at the power-supply traces. In a typical circuit, buck or boost converters generate the power-supply signals. On top of this type of noisy supply, you can receive EFT on top of the power voltages. This EFT will manifest itself as voltage or current spikes. Remember, with higher order geometry designs the circuit works just fine in your current layout. A general rule of thumb going forward is to minimize these effects by managing your power and ground traces (or planes). Finally, your circuit has always required decouple or bypass capacitors, but now the proper selection is critical.

Figure 1 illustrates a few examples of what to do and not do.




Figure 1: Connecting several devices with one ground and VDD trace, (a) can become a candidate for ground and power-supply loops. This topology also enhances power-supply glitches. Having ground or VDD jumper (b) is a better solution but not great. Creating a ground and VDD trace from device to device is a better solution (c) between the three. However, the best solution is to have separate ground and power-supply planes (d) in a multi-layer board.

You have probably heard all of these suggestions before and a multitude of other useful hints beyond this short column. If so, that is a good sign. Since you're a descent engineer, you probably have discovered what you need and don't need in your circuit. I congratulate you. But, change is in the wind. As you move into new silicon geometries, you need to implement those seemingly ancient layout rules. Optimize your layout with power glitches in mind. Add protection circuits such as MOVs and transient suppressors. Protect the I/O, interrupt and reset pins of your controller or processor. Use firmware recovery techniques such as WDT or register refresh so they contain the correct values. All of these techniques will foster the production of a robust design.

Bypass Capacitors: No Black Magic Here

By Bonnie C. Baker, Microchip Technology Inc.

A basic requirement for all electronic circuits is the inclusion of bypass capacitors (also known as decoupling capacitors). These devices reside across the positive supply to ground, as close to the supply pin of the active device as possible. You may get away with excluding these capacitors in low-frequency circuits, but more than not, low-frequency circuits actually have high-frequency entities inside the active devices. An example of a supposed “low-frequency device” is a microcontroller that uses a low-frequency system clock. Granted, the frequency of the clock is slow, but the transition of the internal and I/O gates can occur in nanoseconds. Without proper power-supply filtering, these rising and falling glitches will traverse the circuit. The first step to proper supply filtering is to include a proper-valued, bypass capacitor.

Circuits containing digital devices are not the only systems that require bypass capacitors. Analog devices also benefit from the inclusion of bypass capacitors, but in another way. While bypass capacitors in digital systems control fast rising- and falling-time glitches from the device, bypass capacitors in analog systems assist in preventing power-supply noise from entering the analog device. Typically, analog devices have preventative power-supply filtering built-in, primarily known as power-supply or line-rejection capability. These power-supply rejection (PSR), noise-rejection mechanisms are effective at reducing low-frequency power-supply noise, but this is not the case at higher frequencies.

Once you resign to your fate of including bypass capacitors in your circuit, the re-aiming task is to select the right capacitor value for the various devices in your circuit. Typically, manufacturers will include suggested bypass capacitor values in their data sheets. If the manufacturer does not supply that information, you can determine the proper value on your own.

For instance, with microcontrollers or microprocessors you can calculate the bypass capacitor value when you know the typical rise and fall times of signals from your device (t_{RISE}). You also need the average current while the microcontroller or microprocessor is operating (I_{AVE}). These two quantities are in the microcontroller/microprocessor product data sheet tables. You finally need to define the maximum voltage ripple-noise that you can tolerate on your power-supply trace (V_{RIPPLE}).

Using these values, the appropriate value for the bypass capacitor is:

$$f_{noise} = 1 / (2 * t_{RISE}); \text{ Determine noise frequency}$$

$$I_{SURGE} = I_{AVE} * f_{noise} / f_{micro}; \text{ Approximate surge current}$$

$$C_{BYPASS} = I_{SURGE} / (2 * \pi * f_{noise} * V_{RIPPLE}); \text{ Calculate bypass capacitor value}$$

Analog devices are a benefit in another way. With these kinds of circuits, you need to find the frequency where power-supply noise will affect your circuit. The best place to look to find this information is with the PSR or line-rejection performance over frequency graphs in the product data sheet. Additionally, you need to determine the minimum, acceptable noise that you can tolerate. For instance, with a 12-bit A/D converter you can tolerate un-rejected power-supply noise of approximately $\pm 1/4$ Least Significant Bytes (LSB), for true 12-bit performance. You also need to take a stab at estimating the power-supply, noise-voltage magnitude. With these two parameters, you can refer to the typical PSR versus frequency curve in the manufacturer's data sheet.

For example, Figure 1A provides the PSR over frequency curve of a 12-bit A/D converter. The PSR of this converter is equal to:

$$PSR = 20 \log (V_{POWER-SUPPLY-RIPPLE} / V_{ADC-ERROR})$$

If you determine that the noise level riding on top of your converter power supply is ± 20 mV (or 40 mV peak) and the allowed error is $\pm 1/4$ LSB or 0.61 mV peak (implying 5V full scale range), the noise from the power supply will show up in the output code of the converter as noise as -36.33 dB. This occurs per Figure 1A at approximately 5 MHz. Referring to Figure 1B, the appropriate bypass, ceramic capacitor value for this converter would be between 0.1 μ F and 0.01 μ F.

It is advantageous to note that these calculations use typical values instead of the minimum or maximum numbers. This is important to understand, because selecting the correct bypass capacitor is not an exact science. Not only will devices vary slightly from part to part, the capacitors that you use in your circuit will also vary from part to part as well as over temperature. But, don't let this deter you from using bypass capacitors. The worst of all cases is when you use none.




Figure 1: Using the PSR versus over temperature of an analog part (A) in combination with the ceramic capacitor impedance versus frequency (B) determines the best bypass capacitor for analog parts.

Techniques To Minimize Noise

The Debounce Debacle

By Bonnie C. Baker, Microchip Technology Inc.

Mechanical switches play an important role in many microcontroller and microprocessor applications. These types of switches are simple, inexpensive and easy to install. But, the perceived simplicity of mechanical switches can be deceiving if you intend to keep your user happy. For instance, an annoying, poor switch implementation is a TV button that advances the channel two or three channels higher instead of the intended one-channel advance. More critically, a switch sensor could determine the status of a door. An open door trips an alarm or motivates an attendant to go to the site. False alarms can be disruptive and expensive. Some switches reside in time-critical environments, like in keyboards where a quick, accurate determination of the switch position is critical.

It is easy to blame these problems on the switch, instead of the electronics that sense the switch position. An alternative approach is to examine the dynamic characteristics of your switch and assess its environmental influences. All switches demonstrate a switch contact, “bouncing” action as the switch opens or closes. The switch contacts actually bounce off of each other several times before the contacts settle to their final position. If the switch position is sensitive to the touch, a person who had too much coffee could cause a bouncing action to occur as a result of human touch. Switch manufacturers call this “playing” with the switch. You also need to look at environmental interference, such as vibrations or electromagnetic interference (EMI). Once you’ve evaluated these system dynamics, you are ready to design the electronics in your circuit.

You can implement the switch-interface electronics using an analog or digital solution. The analog solution uses an R/C filter to control the bouncing signal. A digital gate, with hysteresis, takes the output of the R/C filter and sends the switch-position signal into the microcontroller or microprocessor. In this type of system, the bounce time is assumed slower than the R/C time constant. You are probably safe if you design with a 500 ms switch closure time. But, don’t hesitate to verify this timing with your switch on the bench. You will find that switches bounce less as they open as compared to when they close. In this configuration, the microcontroller or microprocessor receives the switch position after the fact.

An alternative to the analog solution is a purely digital one. A digital solution is “almost” free because you no longer need your external components. The analog option has the additional cost of materials, installation and board real estate. The pure digital solution can immediately determine if the switch has changed. The microcontroller or microprocessor can then filter the switch bounce digitally. This debounce circuit can be thought of as a low-pass/decimation filter. The “brute-force debounce” digital solution uses a programmed delay time in firmware to replace the analog R/C time constant.

As an alternative, it is possible to adjust this programmed delay time by using clever programming techniques. This type of algorithm increments a counter when the switch output signal is high and decrements the same counter when the switch output is low. By taking multiple samples, the counter will over range either high or low, which identifies the switch state. The programmer determines the number of and timing between samples. The timing of the samples should not correlate to the main frequency of microcontroller or microprocessor clocks.

Some say the attention to issues surrounding a switch’s open and closure action is a lost art. If this is so, evaluating all the elements of your system is a lost art as well. Electronic systems in the real world usually have a switch somewhere in the circuit. The opening or closing of a switch seldom demonstrates a clean transition. The connection consists of a series of “breaks” and “makes” before the final “make” (or “break”). This bouncing-phenomena can occur several to hundreds of times before things settle. The best, cost-effective solution comes from using the microcontroller/microprocessor approach. You can change your design as you experiment, or on the fly at next to no cost.

References

“My Favorite Software Debouncers”, Jack Ganssle, Embedded Systems Programming Magazine, 6-16-04.

“PICkit™ 1 Flash Starter Kit, User’s Guide”, Microchip Technology Inc., DS40051.

“Examining Switch-Debounce Circuits”, Ron Mancini, EDN, Feb. 21, 2002.

ETP-248: Managing Power, Ground And Noise In Microcontroller/Analog Applications

By Bonnie C. Baker and Keith Curtis, Microchip Technology Inc.

Microcontroller applications often have low-level sensor signals and moderate power-drive circuitry, in addition to the microcontrollers. A peaceful coexistence with these three extremes requires a careful power and ground distribution design. This paper will discuss sources of noise and the paths by which the noise travels. We will cover the theory behind good layout practices and their impact on noise. We will also discuss the proper selection and placement of noise isolating and limiting components. Any designer who struggles to keep digital and power noise out of sensitive input circuits will find this paper useful.

Figure 1 shows a block diagram of the example system that we will use to discuss concepts in this paper. The overall function of this system is to acquire weight and display the results in an LED array, as well as on a laptop computer. Additionally, a fan controller cools the board temperature as needed.

In the example design, there is an analog and digital section. One of the difficulties in this design is the segregation of these two domains. Starting with the analog input to the circuit, the design is capable of measuring weight. The analog interface block in Figure 1 includes a load cell, gain block, anti-aliasing filter and a 12-bit analog-to-digital converter (ADC). The load cell is a wheatstone bridge, as shown in Figure 2.

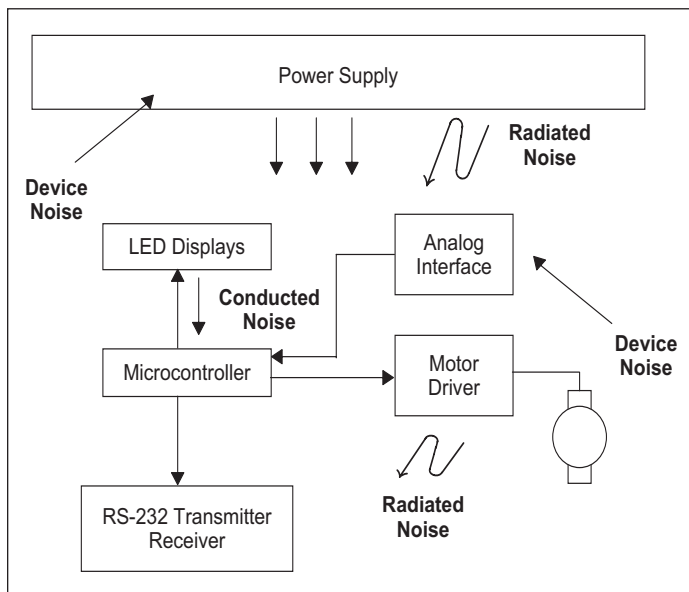


Figure 1: This block diagram models the circuit, along with noise sources, of the system in this paper. The analog interface circuitry measures weight with a load-cell sensor. The interface then transmits those results to a microcontroller. The microcontroller sends the sensor results to an LED display and laptop computer. There is also circuitry for a fan-motor driver.

In the digital section, the microcontroller produces the digital representation of the load-cell value. One of the activities of the microcontroller is to display the measurement results in the LED array. The microcontroller also uses an RS-232 interface port to communicate the data to a desktop computer. The desktop computer takes analog measurements from the microcontroller and displays that data in a histogram chart. Finally, the digital section also includes a PWM driver output for the fan.

This design includes sensitive analog circuitry, a high-power LED display and a potentially noisy digital interface to a laptop computer. The challenge is to design a circuit and layout that allows the coexistence of all of these conflicting elements. We will start by designing the analog section of this circuit and then move on to layout concerns.

Analog Circuit Design

The analog portion of this circuit has a load-cell sensor, dual operational amplifier (MCP6022, Microchip Technology Incorporated) configured as an instrumentation amplifier, a 12-bit 100 kHz SAR ADC (MCP3201, Microchip) and one voltage reference. The ADC's SPI™ port connects directly to a microcontroller (see Figure 2).

The load-cell sensor has a full-scale output range of ± 10 mV. The gain of the instrumentation amplifier (A1 and A2) is 153 V/V. This gain matches the full-scale output swing of the instrumentation amplifier block to the full-scale input range of the ADC. The SAR ADC has an internal input sampling mechanism. With this function, each conversion uses a single sample. The microcontroller acquires the data from the converter and translates the data into a usable format for tasks such as the LED display or the PC interface.

If the implementation of the circuit and layout design of this system is poor (no ground plane, no bypass capacitors and no anti-aliasing filter), it will be an excellent candidate for noise problems. The symptom of a poor implementation is an intolerable level of uncertainty about the digital output results from the ADC. It is easy to assume that this type of symptom indicates that the last device in the signal chain generates the noise problem. But, in fact, the root cause of poor conversion results originates with the PCB layout.

In the worst of circumstances, when noise reduction layout precautions are not taken, the 12-bit system in Figure 2 will output a large distribution of code for a DC input signal. Figure 3 shows data from the output of the converter.

The data in Figure 3 is far from optimum. Six bits of peak-to-peak error changes the 12-bit converter system into a noise-free 9.3-bit system. One can average this data in the digital domain to recapture the full 12-bits, but this will ideally require the averaging of $4^{(12 - 9.3)}$ or 36 samples. In a non-ideal environment, it will be more.

Techniques To Minimize Noise

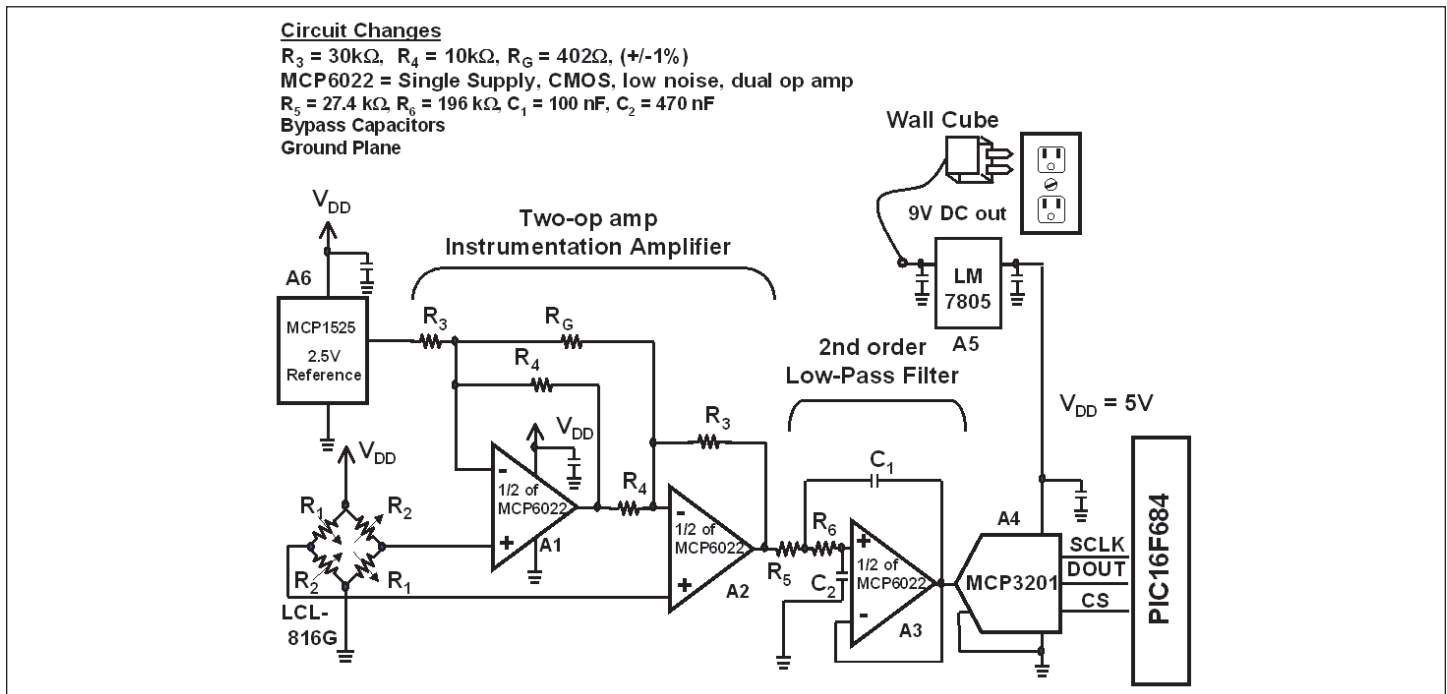


Figure 2: This is the analog portion of the block diagram in Figure 1. The amplifiers and voltage reference tie into the ADC, and the ADC ties to the microcontroller. Building a 2nd order, low-pass filter (A3) requires two resistors, two capacitors and the MCP6022 (Microchip) operational amplifier. The other amplifiers form an instrumentation amplifier circuit. Bypass capacitors are included. For testing, this portion of the circuit is powered from a 110 VAC to 9 VDC wall cube. A linear regulator converts the 9V output of the wall-cube to 5 VDC.

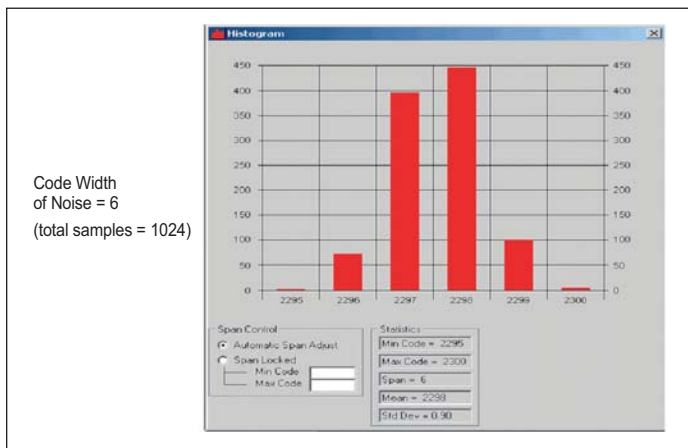


Figure 3: A poor implementation of the 12-bit data acquisition system shown in Figure 2 could easily have an output range of six different codes with a 1024 sample size.

Analog Layout Guidelines

Ground and Power-Supply Strategy: The implementation for the layout of the ground plane is critical when designing low-noise solutions. It is a dangerous practice to omit a ground plane with analog and/or mixed signal devices. Ground planes solve problems such as offset errors, gain errors and noisy circuits. These errors are more prevalent when the layout lacks a ground plane because analog signals typically refer to ground.

When determining the grounding strategy of a board, one should determine if the circuit needs one or more ground planes. If the circuit has a “minimum” amount of digital circuitry onboard, a single ground plane and triple-wide power supply traces may be appropriate. The danger of connecting the digital and analog ground planes together is that the analog circuitry picks-up the noise from the digital return currents. In either case, one should connect the analog grounds, digital grounds and power supplies together at one or more points on the circuit board. The inclusion of a ground plane in a 12-bit system is critical.

ADC Layout Strategies: ADC layout techniques vary with converter technology. When SAR ADCs are used, the entire device should reside on the analog power and ground planes. ADC vendors often supply analog and digital ground pins. If high-resolution SAR converters are used, a digital buffer should be used to isolate the converter from bus activity on the digital side of the circuit. This is the correct strategy to use for delta-sigma ADCs as well.

Figure 4 shows the performance of a board laid out with these considerations. This data shows that the analog section of this circuit works extremely well.

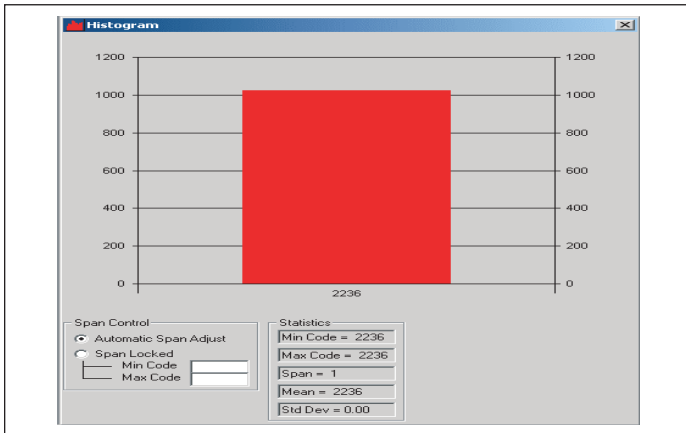


Figure 4: This figure shows the data results from the circuit design in Figure 2. The improvement indicates that our low-noise layout strategies are effective.

Analog Design Conclusion

Concluding this first section, it is important from the beginning of the design to verify that the circuit devices are low noise. In this example, the key elements are the resistors and amplifiers. Following the device selection, make sure that the signal path is properly filtered. This includes the signal path as well as the power-supply train. An uninterrupted ground plane is a critical element in all analog designs. You will eliminate noise that you might otherwise have problems tracking down. Finally, bypass capacitors are critical in analog designs. When you install them, place the capacitors with short leads as close to the power pins as possible.

At this point, the analog design is deemed “good”. However, you will see that the addition of the LED array and fan controller will cause the performance of the analog section to become worse.

Digital Design

We will start with the first-pass digital/analog design integration. In the first pass, we will add the digital section arbitrarily using normal rules of thumb for a layout strategy. This portion of the design adds the LEDs, motor driver, RS-232 transmitter/receiver and microcontroller.

The design uses bypass capacitors and a flyback diode on the motor drive. Bypass capacitors are kept near the IC power, and the ground traces are short. This is done without changing the analog circuit layout. Figure 5 shows the first-pass histogram results from digital/analog layout.

The results of the ADC output from this new board are worse than the first try at the analog section. We are going to recover the original analog behavior by reworking the power and ground strategy. The first corrective action is to separate the digital portions of the power train from the analog portion. Figure 6A shows the first attempt with the analog/digital combination. Figure 6B reflects the second, more successful version.

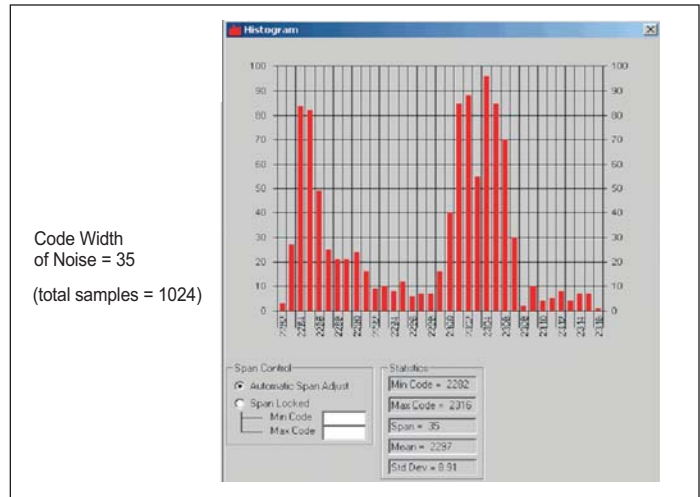


Figure 5: The digital circuit corrupts the analog section, causing a code width of 35 codes from 1024 samples. This data (the output of the ADC) indicates that there are noise sources on the board because of the two features in the data. The results from this data would be difficult to digitally filter to one final ADC code.

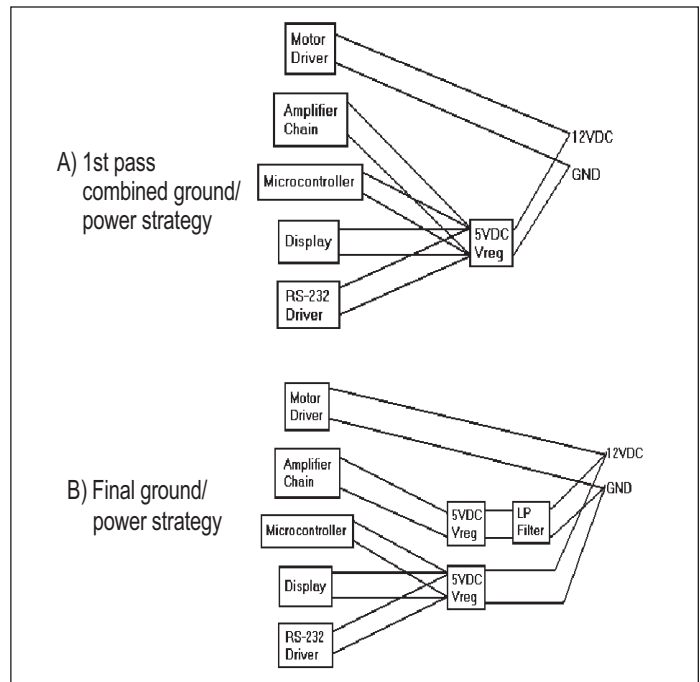


Figure 6: This figure illustrates the first- and second-pass strategies for power and ground. Note that the second pass separates noisy and sensitive circuitry as much as possible.

Techniques To Minimize Noise

The first analog/digital layout routed the digital 5V and ground connections through the analog section. In this configuration, the LED high currents, motor switching and digital controller noise is overlaid onto the sensitive analog power and ground paths (Figure 6A). The pathways for noise in the PCB traces are the power and ground currents interacting with trace resistance and inductance. This produces AC offsets in both the power and ground of the analog portion of the circuit. A quick solution to this problem is to reroute the power and ground traces so that the analog and digital lines are independent, up to a central location. At this central location, they are connected (Figure 6B). This strategy takes advantage of trace resistance, inductance and bypass capacitors to create RC and LC low-pass filters on the power and ground traces. This further isolates noisy and sensitive sections of the design.

Key candidates for radiated noise are the LED traces (that carry high currents), the charge pump in the RS-232 interface (which can pull moderate currents) and the I/Os from the microcontroller (which have fast rise times). The LED and RS-232 driver traces inductively couple noise to adjacent traces in close proximity on the board. This action manifests itself as voltage noise. Fast rise-time signals from the microcontroller capacitively couple into high-impedance, sensitive traces. This coupling activity manifests itself as current noise if traces are in close proximity.

If these factors are considered in the circuit layout, the noise coupled from the noisy digital to the sensitive analog is reduced. The new layout for analog circuitry remains unchanged, as does most of the digital. The difference is that the LED traces now go around the analog block rather than through it. The RS-232 interface's power supply and ground are also kept separate from the sensitive analog and digital functions on the board. And, the power and ground strategy from Figure 6B is used to guide the layout.




Figure 7: Testing the new layout shows that the changes are effective. 1024 samples were collected. The board exhibits true 12-bit operation.

Conclusion

The first step to analog noise reduction is to choose low-noise analog parts. One can remove signal and power noise with filters. Use an anti-aliasing filter as appropriate. In the power bus, use bypass capacitors and inductive chokes as needed. Also, implement the PCB solution with a ground plane.

When you add digital circuitry, develop a ground and power strategy for the entire circuit. In the plan, consider the resistance and inductance of the traces in conjunction with the current density traveling through those paths. The objective in the composite layout is to minimize noise pathways, such as capacitive and inductive coupling between traces, and to use the inductance and resistance of the trace, together with the bypass capacitors, to reduce and isolate noise.

References

- “High-Speed Digital Design: A Handbook of Black Magic”, Howard Johnson, Martin Graham, Prentice Hall, 1993.
- “Noise Reduction Techniques in Electronic Systems”, Henry Ott, John Wiley, N.Y., 1998.
- “The RF Capacitor Handbook”, American Technical Ceramics Inc.
- “The Circuit Designer’s Companion”, Tim Williams.
- “Reference Data for Engineers”, 7th edition, Edward C. Jordan, Editor-in-chief.
- “ABC’s of Transformers & Coils”, Edward J. Bukstein.
- “Techniques that Reduce System Noise in ADC Circuits”, Bonnie C. Baker, ADN007, Microchip Technology Inc.
- “Anti-Aliasing, Analog Filters for Data Acquisition Systems”, Bonnie C. Baker, AN699, Microchip Technology Inc.
- “Interfacing Pressure Sensors to Microchip’s Analog Peripherals”, Bonnie C. Baker, AN695, Microchip Technology Inc.
- “Layout Tips for 12-Bit A/D Converter Applications”, Bonnie C. Baker, AN688, Microchip Technology Inc.
- “Analog Design in a Digital World using Mixed-Signal Controllers”, Keith Curtis, AN823, Microchip Technology Inc.
- FilterLab® Active Filter Design Software, www.microchip.com

Charging High Capacity Lithium Batteries

By George Paparrizos, Microchip Technology Inc.

New portable electronic designs are continuously improving system performance. Designers are adding new features, and the power requirements of these systems are increasing along with their performance. This trend has resulted in the development of Lithium Ion (Li-Ion) battery packs with significantly higher capacities that can provide adequate system run times. Li-Ion batteries have emerged to be the power source of choice for most portable consumer applications. They provide many advantages over other battery types, such as higher energy density and cell voltages. However, as the available battery capacity increases to address today's growing power demands, consumer requirements for shorter battery-charge cycles and smaller physical size are also increasing. Meeting those requirements becomes a major design challenge.

The preferred charging rate for a Li-Ion or Li-Polymer battery pack is 1C, which means that a 1000 mAh battery pack has a preferred fast-charge current of 1A. Charging at this rate provides the shortest charge-cycle times without degradation to the battery-pack performance or its life. To achieve the preferred charging rate while utilizing battery cells with increasingly higher capacity, a higher charge-current level is inevitable. This leads to major design challenges, the most important of which is the heat generated in the charge-management controller IC as well as on the system board. To address these issues, a more careful thermal design is necessary. As shown below, integrating the pass transistor, which dissipates most of the heat, on the charge management IC, becomes, in many cases, a limiting factor for the design performance and operating range.

The power dissipation of a linear-battery charger, the most popular type for 1- and 2-cell Lithium-based applications, can be calculated as follows:

$$P_D = (V_{DD} - V_{BATTERY}) \times I_{CHARGE}$$

With:

P_D : power dissipation

V_{DD} : input (adapter) voltage

$V_{BATTERY}$: battery voltage

I_{CHARGE} : fast charge current

The above formula indicates that a worst-case scenario exists when the input-voltage to battery-voltage differential and the charge current are high. We calculate the power dissipation for a typical scenario, in which the car or wall adapter provides a 5V ± 10 percent supply voltage, the battery voltage is 3.0V and the fast-charge current is 1A. (i.e., 1C charging rate for a 1000 mAh battery) In this case, the power dissipation is calculated to:

$$P_D = (5.5V - 3.0V) \times 1A = 2.5W$$

Table 1 shows the maximum operating ambient temperature assuming a 2.5W power-dissipation scenario, and a maximum junction temperature of 125°C, for some of the most popular package options in the industry.

Package Type	Thermal Resistance (°C/W)*	Max. Ambient Temperature (°C)
MSOP-8	206	-391
SOIC-8	163	-283
MSOP-10	113	-158
SOT-23-6	96	-114
3x3 DFN-8	52	-5
Power SOIC-8	43	18
4x4 QFN-16	38	30
5x6 DFN-8	36	35

Table 1: Max. Operating Ambient Temperature (2.5W Power-Dissipation, Max. Junction Temperature 125°C)

It is apparent from these results that the maximum allowable ambient temperature for most of the package types puts significant limitations on the thermal design. Furthermore, achieving some of these operating conditions is not feasible in a real-world situation. The assumptions for conducting these calculations do not take into account worst-case conditions. For example, the battery might have just been preconditioned (i.e. deeply depleted cells) in which case the battery voltage during the beginning of current regulation could be as low as 2.7V. Additionally, after-market adapters that provide a less tight regulation might be utilized, such that the battery charger would accept input voltages significantly higher than 5V. Lastly, as new applications require more power, battery packs with capacities greater than 1000 mAh are becoming increasingly popular.

Some battery chargers, like Microchip Technology's MCP7384X, provide the ability to use an external P-channel MOSFET. Discrete MOSFETs can run hotter without reliability concerns. They are offered in a variety of heat-dissipating packages that enable higher allowable power dissipation as well as higher maximum operating ambient temperature. The flexibility of selecting a MOSFET that meets the application's power dissipation and electrical requirements allows for higher system performance and cost optimization. Figure 1 demonstrates a typical circuit for this battery-charger family.

The MCP7384X also offers a ± 0.5 percent voltage regulation accuracy (max). Since both over- and under-charging the battery reduces its performance and life, it is essential for the charge management controller IC to provide a high-precision voltage to the battery. Exceeding the termination voltage allows for a short-term gain in available battery energy. However, continuous overcharging can result in a significant reduction in cycle life or safety issues. On the other hand, undercharging reduces the battery cell's available capacity, thereby shortening usable battery time; thus, providing a precise regulation voltage maximizes cell capacity utilization and ensures safe operation. Figures 2 and 3 illustrate that a small decrease in voltage accuracy results in a large decrease in battery capacity.

Charging Lithium Batteries

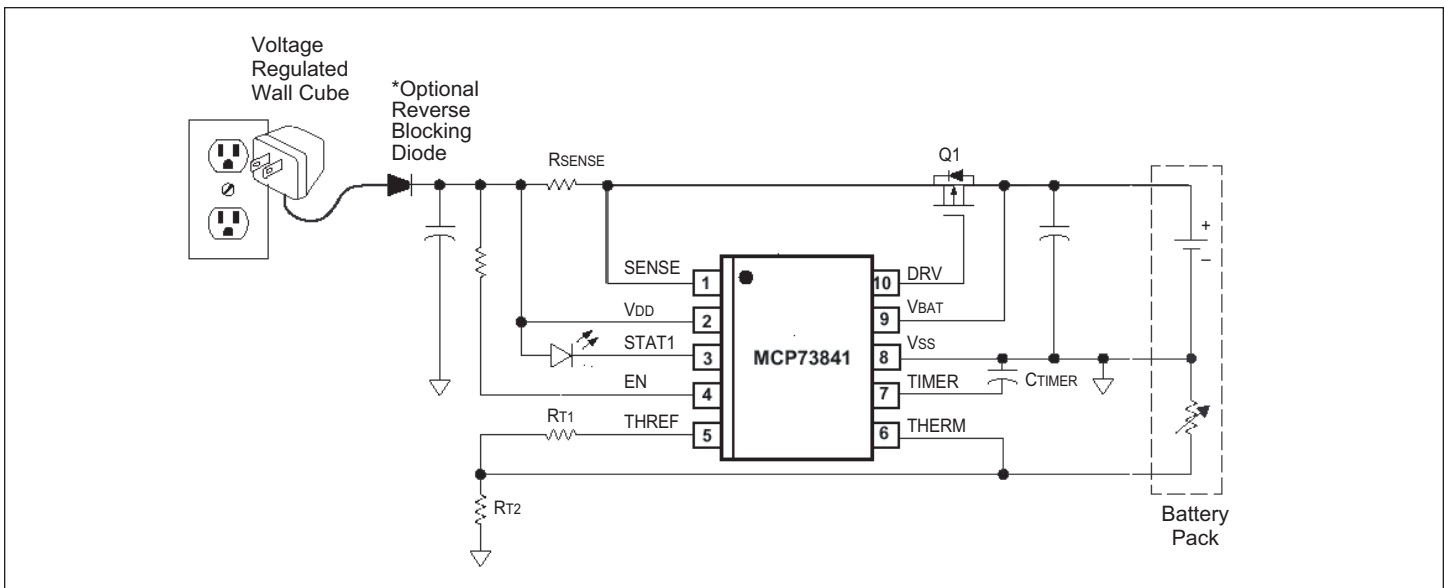


Figure 1: Typical Charging Circuit for High-Capacity Battery Packs.

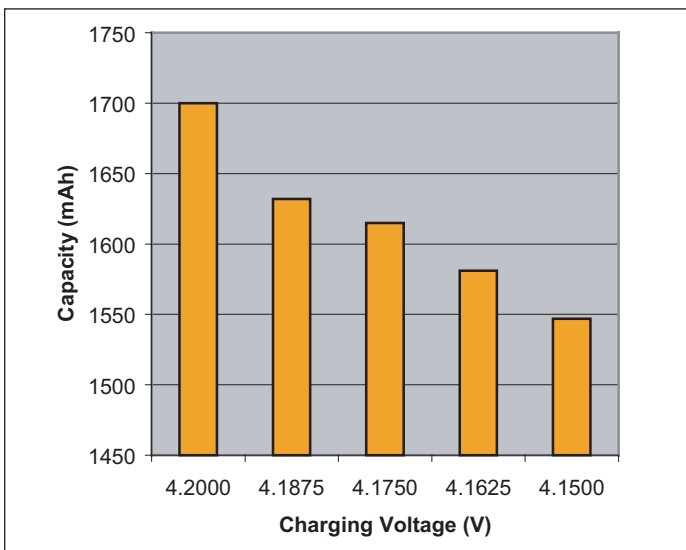


Figure 2: Battery Capacity vs. Charging Voltage.

When the battery voltage falls below a recharge threshold, the MCP7384X products automatically initiate a new charge cycle. This allows the battery to be kept at near-full charge at all times, making it readily available for utilization at any point of time. The operation also allows for higher reliability, since “trickle-charging” Li-Ion batteries for extended periods results in life degradation.

The number of available portable equipment types are expected to increase in the years to come. Extending usable battery life and reducing charge-cycle time are essential for ensuring a positive user experience. Battery manufacturers are already offering new, space-saving Li-Ion battery packs with high capacity to meet today's power demands. While high-capacity batteries are able to extend operating life, they also introduce significant design challenges. Addressing both issues requires accurate voltage regulation and the utilization of external-pass transistors that are capable of providing high charge current while ensuring thermal stability.




Figure 3: Battery Capacity Loss vs. Under-Charge Voltage.

Charging Lithium-Ion Batteries: Not All Charging Systems Are Created Equal

By Scott Dearborn, Microchip Technology Inc.

Introduction

Powering today's portable world poses many challenges for system designers. The use of batteries as a prime power source is on the rise. As a result, a burden has been placed on the system designer to create sophisticated systems utilizing the battery's full potential.

Each application is unique, but one common theme rings through: **maximize battery capacity usage**. This theme directly relates to how energy is properly restored to rechargeable batteries. No single method is ideal for all applications. An understanding of the charging characteristics of the battery and the application's requirements is essential in order to design an appropriate and reliable battery charging system. Each method has its associated advantages and disadvantages. It is the particular application with its individual requirements that determines which method will be the best to use.

Far too often, the charging system is given low priority, especially in cost-sensitive applications. The quality of the charging system, however, plays a key role in the life and reliability of the battery. In this article, the fundamentals of charging Lithium-Ion (Li-Ion) batteries are explored. In particular, linear charging solutions and a microcontroller-based, switch-mode solution shall be explored. Microchip's MCP73843 and MCP73861 linear charge management controllers and PIC16F684 microcontroller along with a MCP1630 pulse width modulator (PWM), shall be used as examples.

Li-Ion Charging

The rate of charge or discharge is often expressed in relation to the capacity of the battery. This rate is known as the C-Rate. The C-Rate equates to a charge or discharge current and is defined as:

$$I = M \times Cn$$

where:

I = charge or discharge current, A

M = multiple or fraction of C

C = numerical value of rated capacity, Ah

n = time in hours at which C is declared.

A battery discharging at a C-rate of 1 will deliver its nominal rated capacity in one hour. For example, if the rated capacity is 1000 mAh, a discharge rate of 1C corresponds to a discharge current of 1000 mA. A rate of C/10 corresponds to a discharge current of 100 mA.

Typically, manufacturers specify the capacity of a battery at a 5 hour rate, n = 5. For example, the above-mentioned battery would provide 5 hours of operating time when discharged at a constant current of 200 mA. In theory, the battery would provide 1 hour of operating time when discharged at a constant current of 1000 mA. In practice, however, the operating time will be less than 1 hour due to inefficiencies in the discharge cycle.

So how is energy properly restored to a Li-Ion battery? The preferred charge algorithm for Li-Ion battery chemistries is a constant, or controlled, current – constant voltage algorithm that can be broken up into four stages: trickle charge, constant current charge, constant voltage charge and charge termination. Refer to Figure 1.




Figure 1: Li-Ion Charge Profile.

Stage 1: Trickle Charge – Trickle charge is employed to restore charge to deeply depleted cells. When the cell voltage is below approximately 3V, the cell is charged with a constant current of 0.1C maximum.

Stage 2: Constant Current Charge – After the cell voltage has risen above the trickle charge threshold, the charge current is raised to perform constant current charging. The constant current charge should be in the 0.2C to 1.0C range. The constant current does not need to be precise and semi-constant current is allowed. Often, in linear chargers, the current is ramped-up as the cell voltage rises in order to minimize heat dissipation in the pass transistor.

Charging at constant current rates above 1C does not reduce the overall charge cycle time and should be avoided. When charging at higher currents, the cell voltage rises more rapidly due to over-voltage in the electrode reactions and the increased voltage across the internal resistance of the cell. The constant current stage becomes shorter, but the overall charge cycle time is not reduced because the percentage of time in the constant voltage stage increases proportionately.

Stage 3: Constant Voltage – Constant current charge ends and the constant voltage stage is invoked when the cell voltage reaches 4.2V. In order to maximize performance, the voltage regulation tolerance should be better than +1%.

Stage 4: Charge Termination – Unlike nickel-based batteries, it is not recommended to continue to trickle charge Li-Ion batteries. Continuing to trickle charge can cause plating of metallic lithium, a condition that makes the battery unstable. The result can be sudden, automatic and rapid disassembly.

Charging is typically terminated by one of two methods: minimum charge current or a timer (or a combination of the two). The minimum current approach monitors the charge current during the constant voltage stage and terminates the charge when the charge current diminishes in the range of 0.02C to 0.07C. The second method determines when the constant voltage stage is invoked. Charging continues for an additional two hours, and then the charge is terminated.

Charging Lithium Batteries

Charging in this manner replenishes a deeply depleted battery in roughly 2.5 to 3 hours.

Advanced chargers employ additional safety features. For example the charge is suspended if the cell temperature is outside a specified window, typically 0°C to 45°C.

Li-ion Charging – System Considerations

A high-performance charging system is required to recharge any battery quickly and reliably. The following system parameters should be considered in order to ensure a reliable, cost-effective solution.

Input Source

Many applications use very inexpensive wall cubes for the input supply. The output voltage is highly dependent on the AC input voltage and the load current being drawn from the wall cube.

In the US, the AC mains input voltage can vary from 90 VRMS to 132 VRMS for a standard wall outlet. Assuming a nominal input voltage of 120 VRMS, the tolerance is +10%, -25%. The charger must provide proper regulation to the battery independent of its input voltage. The input voltage to the charger will scale in accordance to the AC mains voltage and the charge current:

$$V_o = \sqrt{2} \times V_{IN} \times a - 1o (R_{EQ} + R_{PTC}) - 2 \times V_{FD}$$

R_{EQ} is the resistance of the secondary winding plus the reflected resistance of the primary winding (RP/a^2).

R_{PTC} is the resistance of the PTC, and V_{FD} is the forward drop of the bridge rectifiers. In addition, transformer core loss will slightly reduce the output voltage.

Applications that charge from a car adapter can experience a similar problem. The output voltage of car adapter will have a typical range of 9V to 18V.

Constant Current Charge Rate and Accuracy

The choice of topology for a given application may be determined by the desired constant current. Many high constant current, or multiple cell applications rely on a switch-mode charging solution for improved efficiency and less heat generation.

Linear solutions are desirable in low to moderate fast charge current applications for their superior size and cost considerations. However, a linear solution purposely dissipates excess power in the form of heat.

The tolerance on the constant current charge becomes extremely important to a linear system. If the regulation tolerance is loose, pass transistors and other components will need to be oversized adding size and cost. In addition, if the constant current charge is low, the complete charge cycle will be extended.

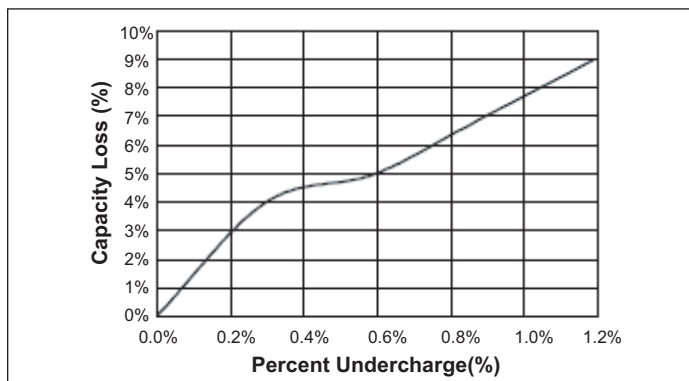


Figure 2: Capacity Loss vs. Undercharge Voltage.

Output Voltage Regulation Accuracy

The output voltage regulation accuracy is critical in order to obtain the desired goal: maximize battery capacity usage. A small decrease in output voltage accuracy results in a large decrease in capacity. However, the output voltage can not be set arbitrarily high because of safety and reliability concerns.

Figure 2 depicts the importance of output voltage regulation accuracy.

Charge Termination Method

It can not be stressed enough that over charging is the Achilles' heel of Li-Ion cells. Accurate charge termination methods are essential for a safe, reliable, charging system.

Cell Temperature Monitoring

The temperature range over which a Li-Ion battery should be charged is 0°C to 45°C, typically. Charging the battery at temperatures outside of this range may cause the battery to become hot. During a charge cycle, the pressure inside the battery increases causing the battery to swell. Temperature and pressure are directly related. As the temperature rises, the pressure can become excessive. This can lead to a mechanical breakdown inside the battery or venting. Charging the battery outside of this temperature range may also harm the performance of the battery or reduce the battery's life expectancy.

Generally, thermistors are included in Lithium-Ion battery packs in order to accurately measure the battery temperature. The charger measures the resistance value of the thermistor between the thermistor terminal and the negative terminal. Charging is inhibited when the resistance, and therefore the temperature, is outside the specified operating range.

Battery Discharge Current Or Reverse Leakage Current

In many applications, the charging system remains connected to the battery in the absence of input power. The charging system should minimize the current drain from the battery when input power is not present. The maximum current drain should be below a few microamperes and, typically, should be below one microampere.

Li-ion Charging – Application Examples

Taking the above system considerations into account, an appropriate charge management system can be developed.

Linear Solutions

Linear charging solutions are generally employed when a well-regulated input source is available. Linear solutions, in these applications, offer advantages of ease of use, size and cost.

Due to the low efficiency of a linear charging solution, the most important factor is the thermal design. The thermal design is a direct function of the input voltage, charge current and thermal impedance between the pass transistor and the ambient cooling air. The worst-case situation is when the device transitions from the trickle charge stage to the constant current stage. In this situation, the pass transistor has to dissipate the maximum power. A trade-off must be made between the charge current, size, cost and thermal requirements of the charging system.

Take, for example, an application required to charge a 1000 mAh, single Li-Ion cell from a 5V \pm 5% input at a constant current charge rate of 0.5C or 1C. Figure 3 depicts Microchip's MCP73843 used to produce a low cost, stand-alone solution. With a few external components, the preferred charge algorithm is implemented.

The MCP73843 combines high accuracy constant current, constant voltage regulation with automatic charge termination.

Charging Lithium Batteries

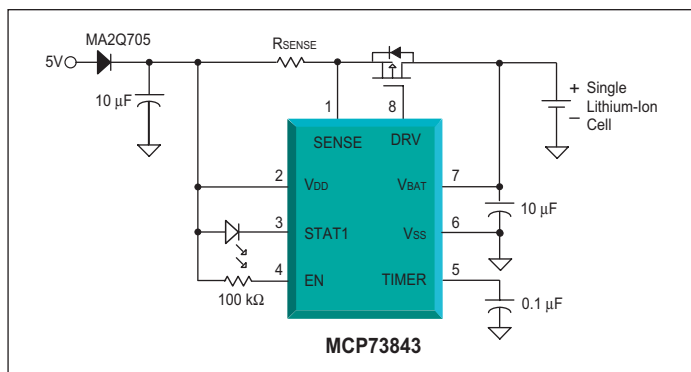


Figure 3: Typical Linear Solution.

In an effort to further reduce size, cost and complexity of linear solutions, many of the external components can be integrated into the charge management controller. Advanced packaging and reduced flexibility come along with higher integration. These packages require advanced equipment for manufacturing, and, in many instances, preclude rework. Typically, integration encompasses charge current sensing, the pass transistor and reverse discharge protection. In addition, these charge management controllers typically employ some type of thermal regulation. Thermal regulation optimizes the charge cycle time while maintaining device reliability by limiting the charge current based on the device die temperature. Thermal regulation greatly reduces the thermal design effort.

Figure 4 depicts a fully integrated, linear solution utilizing Microchip's MCP73861. The MCP73861 incorporates all the features of the MCP73843 along with charge current sensing, the pass transistor, reverse discharge protection and cell temperature monitoring.

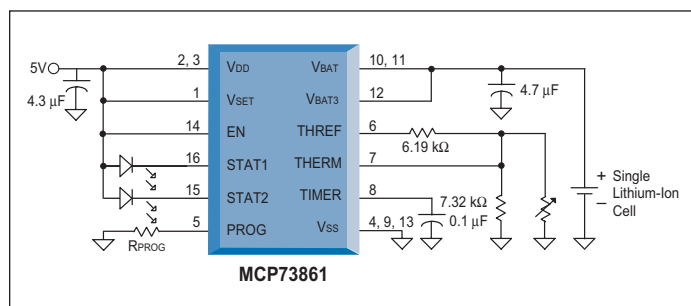


Figure 4: Typical, Fully Integrated, Linear Solution.

Charge Cycle Waveforms

Figure 5 depicts complete charge cycles utilizing the MCP73843 with constant current charge rates of 1C and 0.5C. Charging at a rate of 0.5C instead of 1C, it takes about 1 hour longer for the end of charge to be reached. The MCP73843 scales the charge termination current proportionately with the fast charge current. The result is an increase of 36% in charge time with the benefit of a 2% gain in capacity and reduced power dissipation. The change in termination current from 0.07C to 0.035C results in an increase in final capacity from ~98% to ~100%. The system designer has to make a trade-off between charge time, power dissipation and available capacity.




Figure 5: MCP73843 Charge Cycle Waveforms.

Switch Mode Charging Solutions

Switch mode charging solutions are generally employed in applications that have a wide ranging input or a high input to output voltage differential. In these applications, switch mode solutions have the advantage of improved efficiency. The disadvantage is system complexity, size and cost.

Take, for example, an application required to charge a 2200 mAh, single Li-Ion cell from a car adapter at a constant current charge rate of 0.5C or 1C. It would be extremely difficult to utilize a linear solution in this application due to the thermal issues involved. A linear solution employing thermal regulation could be utilized, but the charge cycle times at the reduced charge currents may be prohibitive.

The first step to designing a successful switch mode charging solution is to choose a topology: buck, boost, buck-boost, flyback, Single-Ended Primary Inductive Converter (SEPIC) or other. Knowing the input and output requirements, and experience, quickly narrows the choices down to two for this application: buck or SEPIC. A buck converter has the advantage of requiring a single inductor. Disadvantages of this topology include an additional diode required for reverse discharge protection, high-side gate drive and current sense and pulsed input current (EMI concern).

The SEPIC topology has advantages that include lowside gate drive and current sense, continuous input current and dc isolation from input to output. The main disadvantage of the SEPIC topology is the use of two inductors and an energy transfer capacitor.

Figure 6 depicts a schematic for a switch mode charger. Microchip's high speed Pulse Width Modulator (PWM), MCP1630, has been utilized in a pseudo smart battery charger application. The MCP1630 is a high-speed, microcontroller adaptable, pulse width modulator. When used in conjunction with a microcontroller, the MCP1630 will control the power system duty cycle to provide output voltage or current regulation. The microcontroller, PIC16F684, can be used to regulate output voltage or current, switching frequency and maximum duty cycle. The MCP1630 generates duty cycle, and provides fast over current protection based off various external inputs. External signals include the input oscillator, the reference voltage, the feedback voltage and the current sense. The output signal is a square-wave pulse. The power train used for the charger is SEPIC.

The microcontroller provides an enormous amount of design flexibility. In addition, the microcontroller can communicate with a battery monitor (Microchip's PS700) inside the battery pack to significantly reduce charge cycle times.

Charging Lithium Batteries

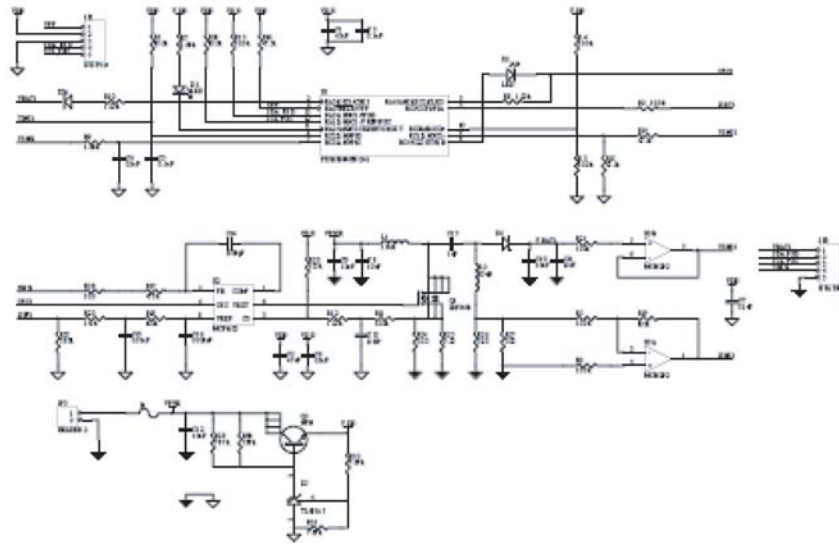


Figure 6: Switch Mode SEPIC Charger.

Charge Cycle Waveforms

Figure 7 depicts complete charge cycles utilizing the switch mode charging solution. By utilizing a battery monitor in the charging system, charge cycles can be significantly reduced. The battery monitor eliminates sensing the voltage produced across the packs protection circuitry and contact resistance by the charging current.




Figure 7: Charge Cycles Using a Switch Mode Charging Solution.

Conclusion

Properly restoring energy using the latest battery technology for today's portable products requires careful consideration. An understanding of the charging characteristics of the battery and the application's requirements is essential in order to design an appropriate and reliable battery charging system.

Linear and switch mode charging solutions for Li-Ion batteries were presented. The guidelines and considerations presented herein should be taken into account when developing any battery charging system.

References

- "Handbook of Batteries Third Edition", David Linden, Thomas B. Reddy, New York: McGraw-Hill, Inc., 2002.
- http://sanyo.com/batteries/lithium_ion.cfm
- "MCP73843 Data Sheet", Microchip Technology Inc., DS21823.
- "MCP73861 Data Sheet", Microchip Technology Inc., DS21893.
- "MCP1630 Data Sheet", Microchip Technology Inc., DS21896.
- "PIC16F684 Data Sheet", Microchip Technology Inc., DS41202.
- "PS700 Data Sheet", Microchip Technology Inc., DS21760.

Working Your Amplifier Inside The Single-Supply Voltage “Box”

By Bonnie C. Baker, Microchip Technology Inc.

It may seem easy enough to transfer classical dual-supply, operational-amplifier circuits directly into a single-supply environment. This is true for a few circuits, but now your amplifier output is swinging much closer to the supply rails than before, and your ground reference has disappeared. In this article, we are going to examine standard voltage-feedback, amplifier circuits (such as non-inverting gain) inverting gain, difference amplifiers, instrumentation amplifiers and a photo-sensing configuration. Within these discussions, we will explore the advantages and limitations of working inside the single-supply environment.

Issues to Consider When You Convert from Dual Supply to a Single Supply

You could say that an operational amplifier is an operational amplifier, regardless of the supply voltage that is used. The general characteristics from one device to another are consistent. All voltage-feedback amplifiers have two high-impedance inputs and a low-impedance output. When in a closed-loop circuit, the voltage of the two inputs of the amplifier track each other. The amplifier open-loop gain from input to output is usually above 80 dB or 10,000 V/V. These general characteristics are not the issues that get you into trouble when you convert from dual-supplies to a single-supply. In this article, the voltage of the dual supplies is $\pm 15\text{V}$ and the voltage of the single supply is 0-5V.

You should focus on two performance characteristics when performing this conversion. These two issues are the input-voltage range and output-voltage swing. The input and output characteristics are fully specified in single-supply operational amplifier data sheets. Generally, if you are converting to a single-supply amplifier circuit these specifications, from amplifier to amplifier, are very close. However, if you don't account for these performance limitations, you will drive the input or output of your amplifier well outside the operating ranges. Outside these ranges the “good” and the “best” single-supply amplifiers act the same.

In this discussion, it is presumed that you understand the circuit design topologies of single-supply amplifiers. “Operational Amplifiers Part 1 of 6: What Does “Rail-to-Rail” Input Operation Really Mean?”⁽¹⁾ discusses the details of the topology of the single-supply amplifier input stages. If you read this article, you will find that when common-mode voltage goes beyond the capabilities of the input transistors the output of the amplifier will latch to either rail. “Operational Amplifiers Part 2 of 6: Working with Single-Supply Operational-Amplifier Output Characteristics”⁽²⁾ discusses the details of single-supply amplifier output-stage performance. This article defines two areas of the amplifier output-stage operation. If you drive your amplifier “hard” into either rail (within a few 10s of millivolts), the amplifier will leave its linear region. If you only drive your amplifier a few hundred millivolts from the rails (as defined in the open-loop gain specification), the amplifier will perform in its linear region.

We will concentrate on these two performance characteristics. You will find that this gives you enough guidance to successfully convert all of your amplifier dual power-supply circuits to single supply. And, as stated before, besides these two areas, an amplifier is an amplifier.

The Buffer Amplifier Configuration

The buffer is the easiest amplifier circuit to transfer from a dual-supply to a single-supply environment. Figure 1 shows an amplifier configured as a buffer.




Figure 1: The amplifier buffer operates with the same circuit topology for a dual or single-supply application.

When an amplifier is configured using a single-supply source, there are a few unexpected limitations of the input stage and output stage. The input stage of single-supply amplifiers can stop the signal from entering the amplifier. If the data sheet of the amplifier does not claim rail-to-rail operation, it most likely will not be a rail-to-rail input amplifier. This limits your input swing. If the input to the amplifier, V_{IN} , goes above the limits of the input transistor, the output travels to the positive rail.

The output stage also limits the performance of this circuit in a single-supply configuration. With dual-supplies, the output stage is able to perform across its full output range without distortion. For instance, if the amplifier in Figure 1 had a $\pm 15\text{V}$ supply and the input, V_{IN} , was 10V, V_{OUT} would be 10V also (assuming no offset errors exist). This conclusion is obvious. However, if you use 0-5V supplies on the same circuit, with an input voltage of zero volts the output will not produce the same voltage. The output will ride 10s of millivolts above ground. The actual limit for this value is the low-level, output-voltage swing and dependant on the particular amplifier you are using.

If you are counting on sending a negative voltage through this circuit, it will not work. Again, this may seem obvious. However, you are converting from a dual-supply to a single-supply environment. If the electronics before the amplifier still use a dual supply, this can be a problem.

The phenomena will also happen at the positive, output-voltage rail. If you are using a rail-to-rail input amplifier and you drive the output high, the output will fall short of reaching the supply rail, or 5V. It will actually fall short by 10s of millivolts below the supply. The amplifier's data sheet calls out this actual limit as the high-level output-voltage swing.

This may not seem like a significant problem unless you are counting on using these extreme voltages in your system. For instance, if your amplifier is driving an A/D converter that has an input range of 0 to 5V, several codes on the bottom and top of your digital output word will never appear.

Operational Amplifiers

Non-Inverting Configurations With Gain Built-in

The non-inverting configuration is more forgiving than the buffer circuit. With the non-inverting amplifier configuration (Figure 2.), the gain of the circuit is greater than +1V/V. This circuit will send the voltage applied to V_{IN} and gain it using the resistors in the circuit.




Figure 2: The non-inverting, gained amplifier circuit is also easy to transfer from a dual-supply environment to a single-supply system. Since the signal is gained by the equation $(1 + R_2/R_1)$ to the output, the input stage is not exercised as much as it is with the buffer. The output stage is more likely to over-range because of the gain in the system.

This circuit is not hard to convert between a dual- and single-supply environment. The input stage limitations do not come into play as much as they did with the buffer circuit. Most single-supply amplifiers have an input, common-mode-range to ground. This amplifier circuit will gain any signal on the input, V_{IN} , to a higher value. It is probable that the output stage will limit the signal on the positive supply rail before the input stage can reach any internal limits.

Photo-Sensing Problems

The transimpedance amplifier that is used for photo-sensing applications is a little more difficult to convert from dual- to single-supply environments. A primary issue is to correctly bias the photo sensor across the input terminals. Figure 3 illustrates the proper connection.




Figure 3: Grounding the anode of the photo sensor and attaching cathode to R_F allows the sensor current to flow through R_F creating a positive voltage at V_{OUT} . If the polarity of the photo sensor is reversed across the amplifier-input terminals, the sensor current will be reversed causing the output of the amplifier to remain on the ground rail.

In this configuration (Figure 3), the current from the photo diode during excitation from a light source will change the output of the amplifier in a positive direction. This arrangement appears to work well except for the fact that the amplifier's output is unable to swing all the way down to ground. Consequently, smaller magnitudes of luminance from the light source will be not be registered by the A/D converter.

Figure 4 provides a solution to this problem. The output of the amplifier is raised by 300 mV (from ground) to ensure that the amplifier is operating in its linear region. A level-shift network, consisting of R_1 , R_2 and A_2 , achieves this goal.




Figure 4: R_1 and R_2 , in conjunction with V_{DD} , form a voltage divider. The output of this divider is 300 mV. A buffer, A_2 , isolates the impedance of R_1 and R_2 while connecting to the anode of the photo sensor. This circuit configuration creates a 300 mV level shift of the signal to the output of the transimpedance amplifier (A_1). The A/D converter also has this level shift applied to its inputs.

The conditions of the open-loop gain specification of the MCP6022 are " $V_{OUT} = V_{SS} + 300 \text{ mV}$ and $V_{DD} - 300 \text{ mV}$." This circuit creates a 300 mV level shift of the signal so that the amplifier, A_1 , remains in its linear region. The selection of A_2 is critical. A_2 must be able to supply current to the photo-sensing circuit in a timely manner. Consequently, A_2 must be as fast or faster than A_1 . Dual amplifiers work well with these requirements. A precision voltage reference (replacing V_{DD}) at the top of R_1 will add stability and reduce noise.

The Inverting Configuration Will Surprise You

The inverting-amplifier configuration (Figure 5) will only work in a single-supply circuit if you have a voltage reference. You can imagine that by connecting the voltage reference node, V_{SHIFT} , to ground, the circuit would only work if the input signal were negative. Figure 6 shows one example of a voltage-reference circuit.

Operational Amplifiers




Figure 5: In a single-supply circuit, you need a voltage to make this system work correctly. By connecting V_{SHIFT} pin to ground, the output will attempt to go below ground with a positive input. This circuit will work well with a V_{SHIFT} ground connect if the input voltage is negative. But, you have to ask yourself, “where will that negative voltage come from in a single-supply circuit?”




Figure 6: Using this type of voltage reference circuit to drive the V_{SHIFT} (Figure 5) input of an amplifier configured in an inverting gain will level shift the output of the circuit amplifier into its linear region. Try to set the V_{SHIFT} voltage around the center of the input range of the circuit.

The Difference Amplifier

Figure 7 illustrates an implementation of the difference-amplifier function. The DC transfer function of this circuit is equal to:

$$V_{OUT} = V_1 \frac{R_4(R_1+R_2)}{((R_3+R_4)R_1)} - V_2 \frac{(R_2/R_1)}{R_3 \frac{(R_1+R_2)}{((R_3+R_4)R_1)}}$$

If R_1/R_2 is equal to R_3/R_4 , the closed-loop output of this circuit equals:

$$V_{OUT} = (V_1 - V_2)(R_2/R_1) + V_{SHIFT}$$

In this circuit, the maximum allowable voltage of the amplifier's inputs limit the input common-mode range. It is possible to have input voltages at V_1 and V_2 that exceed the power-supply voltages. The gain on both input signals is equal. This resistor configuration also subtracts the common-mode voltage of the two input signals. In addition, by setting the two resistor ratios to be greater than one, you can easily implement a gain larger than one.

In a single-supply environment, a voltage reference (V_{SHIFT}) can center the output signal between ground and the power supply. Otherwise, it is possible to drive the output beyond the ground or V_{DD} rail. The purpose and effects of this reference voltage is to shift the output signal into the linear region of the amplifier.




Figure 7: The difference amplifier operates best in a dual-supply environment. In a dual-supply environment, connect V_{Dual-} and single-supply can power the two amp-instrumentation amplifier. Single-supply circuits require a center-supply reference, V_{SHIFT} to ground. In the single-supply environment, where there is no “mid-supply” connect, connecting the V_{SHIFT} pin to $\frac{1}{2} V_{DD}$ is appropriate.

Figure 8 illustrates two possible circuits for the voltage, V_{SHIFT} .




Figure 8: A precision voltage reference (A) or a resistor divider (B) between V_{DD} and ground is used to generate the voltage at V_{SHIFT} in Figure 7. Single-supply difference amplifiers, three-amp instrumentation amplifiers and the two-amp instrumentation amplifiers use these circuits.

The precision, voltage-reference device in Figure 8A is a suitable high-precision solution for your single-supply circuits. This voltage reference will give you accurate, DC results at room and over temperature. Figure 8B illustrates an alternative solution. In this circuit, splitting the resistor, R_4 (Figure 7) between V_{DD} and ground provides a voltage for V_{SHIFT} . Figure 8 summarizes the governing equations for the voltage of V_{SHIFT} and gain of the difference amplifier. The accuracy of this circuit depends on the resistor matching and stability of V_{DD} .

Operational Amplifiers

The Three-Amp Instrumentation Amplifier

The most versatile instrumentation amplifier configuration uses three operational amplifiers in its implementation. This instrumentation amplifier is easy to understand because each of the three operational amplifiers serves a specific function. With this circuit configuration (shown in Figure 9), two of the three operational amplifiers (A1 and A2) gain the two input signals. The third amplifier, A3, subtracts the two gained input signals, thereby providing a single-ended output. The transfer function of the circuit is equal to:

$$V_{OUT} = (V_{IN+} - V_{IN-})(1 + 2R_F/R_G)(R_2/R_1)$$

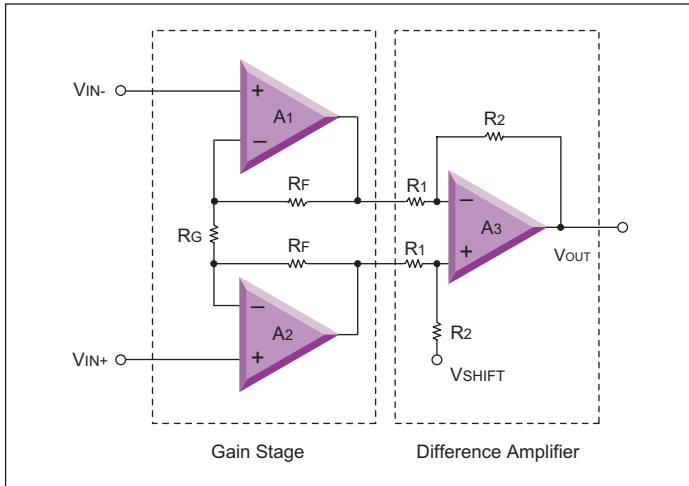


Figure 9: The three-amp instrumentation amplifier has two fundamental stages. The first stage (gain stage) gains the input signals and presents the results to the second stage (difference amplifier).

In single-supply applications, the circuits in Figure 8 generate the center-supply reference, V_{SHIFT} .

The Two-Amp Instrumentation Amplifier

The design shown in Figure 10 uses two operational amplifiers. This design configuration is typically called the two op-amp instrumentation amplifier.




Figure 10: Dual- and single-supply can power the two amp instrumentation amplifier. Single-supply circuits require a center-supply reference, V_{SHIFT} .

Discrete designs use dual amplifiers for good matching of bandwidth and performance over temperature. This instrumentation amplifier utilizes the high impedance of the non-inverting input of the operational amplifiers, thereby significantly reducing source impedance mismatch problems at DC. The transfer function of this circuit is equal to:

$$V_{OUT} = (V_{IN+} - V_{IN-})(1 + R_1/R_2 + 2R_1/R_G) + V_{SHIFT}$$

If the application is in a single-supply environment, this circuit will typically require a reference that is half of the way between the power-supply voltages. In Figure 10, V_{SHIFT} serves that function. This circuit does not allow for zero-volt, common mode, input voltages in single-supply systems.

The center-supply reference, V_{SHIFT} , is implemented using circuits in Figure 8.

Conclusion

The task of transferring your dual-supply amplifier circuits to single-supply environments is straightforward. During the conversion, pay attention to input common-mode range and output-swing specification violations. If any violations of the single-supply amplifier's input or output stage occur as a result of the new power supplies, the addition of a voltage reference can eliminate these violations, making the circuit usable.

References

"Operational Amplifiers Part 1 of 6: What Does "Rail-to-Rail" Input Operation Really Mean?", Bonnie C. Baker, AnalogZone, March 22, 2004.

"Operational Amplifiers Part 2 of 6: Working with Single-Supply Operational-Amplifier Output Characteristics", Bonnie C. Baker, AnalogZone, May 3, 2004.

Every Amplifier Is Waiting To Oscillate And Every Oscillator Is Waiting To Amplify

By Bonnie C. Baker, Microchip Technology Inc.

What is operational amplifier (op amp) circuit stability, and how do you know when you are on the “hairy edge”? Typically, there is a feedback system around the op amp to stabilize the variability and reduce the magnitude of the open-loop gain from part to part. With this approach, the stability of your amplifier circuit depends on the variability of the resistors in your circuit, not your op amps. Using resistors around your op amp provides circuit “stability”. At least you hope that a predictable gain is ensured. But, it is possible to design an amplifier circuit that does quite the opposite.

You can design an amplifier circuit that is extremely unstable to the point of oscillation. In these circuits, the closed gain is somewhat trivial, because an oscillation is “swamping out” your results at the output of the amplifier. In a closed-loop amplifier system, stability can be ensured if you know the phase margin of the amplifier system. In this evaluation, the Bode stability-analysis technique is commonly used. With this technique, the magnitude (in dB) and phase response (in degrees) of both the open-loop response of the amplifier and circuit-feedback factor are included in the Bode plot. This article looks at these concepts and makes suggestions on how to avoid the design of a “singing” circuit when your primary feedback is frequency stability.

The Internal Basics of the Operational-Amplifier Block Diagram

Before getting started on the frequency analysis of an amplifier circuit, let's review a few amplifier topology concepts. Figure 1 shows the critical internal op-amp elements that you need to be familiar with if you engage in a frequency analysis. This amplifier has five terminals, as expected, but it also has parasitics, such as input capacitance and frequency dependent open-loop gain.

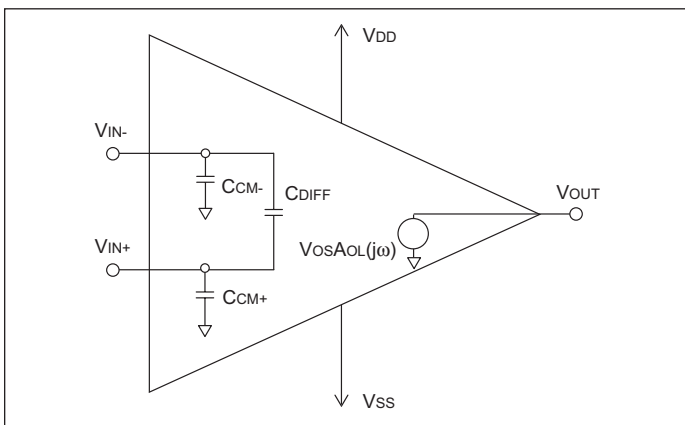


Figure 1: The voltage-feedback operational amplifier frequency model includes the input capacitances (C_{DIFF} and C_{CM}) to ensure that the interaction of the external input-source parasitics and the feedback parasitics can be taken into account in a frequency evaluation. This model also has the internal open-loop gain over frequency ($A_{OL}(j\omega)$). These two parameters ensure that the internal parasitics of the output stage are in the analysis.

The two input terminals have a common-mode capacitance to ground (C_{CM}) and differential capacitance (C_{DIFF}) between the inputs. $A_{OL}(j\omega)$ represents the open-loop gain, frequency response of the amplifier. Figure 2 shows the frequency response of a typical voltage-feedback amplifier using the Bode-plot method.




Figure 2: A gain plot and a phase plot illustrate the frequency behavior of a voltage-feedback amplifier. In this simple representation, the amplifier has two dominant poles. The first pole occurs at lower frequencies, typically between 10 Hz to 1 kHz (depending on the gain-bandwidth product of the amplifier). The second pole resides at higher frequencies. This pole occurs at a higher frequency than the zero decibels (dB) frequency. If it is lower, the amplifier is usually unstable in a unity-gain circuit.

The amplifier gain response over frequency ($A_{OL}(j\omega)$) for a voltage-feedback amplifier is usually modeled with a simple second order transfer function. This second-order transfer function has two poles. The two plots in Figure 2 illustrate the gain (top) and phase response (bottom) of a typical op amp. The units of the y-axis of the gain curve are dB.

Ideally, the open-loop gain of an amplifier is equal to the magnitude of the ratio of the voltage at the output terminal, divided by the difference of the voltages applied between the two input terminals.

$$A_{OL}(\text{dB}) = 20 \log \left| \left(V_{OUT} / (V_{IN+} - V_{IN-}) \right) \right|$$

It would be nice if this open-loop gain ratio were infinite. But in reality, the complete frequency response of the open-loop gain, $A_{OL}(j\omega)$, is less than ideal at DC, FT attenuates at a rate of 20 dB/decade. This attenuation starts at the frequency where the first pole in the transfer function appears. This is illustrated in the Bode plot in Figure 2.

Usually, the first pole of the open-loop response of an operational amplifier occurs between 1 Hz to 1 kHz. The second pole occurs at a higher frequency, nearer to where the open-loop gain-curve crosses 0 dB. The gain response of an amplifier starts to fall off at 40 dB/decade, which is at the frequency where the second pole occurs.

Operational Amplifiers

The units of the y-axis of the phase plot in Figure 2 (bottom plot) are degrees. You can convert degrees to radians with the following formula:

$$\text{Phase in radians} = (\text{Phase in degrees}) * 2\pi / 360^\circ$$

Phase in degrees can be translated to phase delay or group delay (seconds) with the following formula:

$$\text{Phase delay} = -(\delta\text{phase}/\delta f)/360^\circ$$

The same x-axis frequency scale aligns across both plots.

The phase response of an amplifier in this open-loop configuration is also predictable. The phase shift or change from the non-inverting input to the output of the amplifier is zero degrees at DC. Conversely, the phase shift from the inverting input terminal to the output is equal to -180 degrees at DC.

At one decade ($1/10 f_1$) before the first pole, f_1 , the phase relationship of non-inverting input to output has already started to change, ~ -5.7 degrees. At the frequency where the first pole appears in the open-loop gain curve (f_1), the phase margin has dropped to -45 degrees. The phase continues to drop for another decade ($10f_1$) where it is 5.7 degrees above its final value of -90 degrees. These phase-response changes are repeated for the second pole, f_2 .

What is important to understand are the ramifications of changes in this phase relationship with respect to the input and output of the amplifier. One frequency decade past the second pole, the phase shift of the non-inverting input to output is ~ -180 degrees. At this same frequency, the phase shift of the inverting input to output is zero or ~ -360 degrees. With this type of shift, V_{IN+} is actually inverting the signal to the output. In other words, the roles of the two inputs have reversed. If the role of either of the inputs changes like this, the amplifier will ring as the signal goes from the input to the output in a closed-loop system. The only thing stopping this condition from occurring with the stand-alone amplifier is that the gain drops below 0 dB. If the open-loop gain of the amplifier drops below 0 dB in a closed-loop system, the feedback is essentially “turned OFF”.

Stability in Closed-Loop Amplifier Systems

Typically, op amps have a feedback network around them. This reduces the variability of the open-loop gain response from part to part. Figure 3 shows a block diagram of this type of network.




Figure 3: A block diagram of an amplifier circuit, which includes the amplifier-gain cell, A_{OL} and the feedback network, B .

In Figure 3, $\beta(j\omega)$ represents the feedback factor. Due to the fact that the open-loop gain of the amplifier (A_{OL}) is relatively large and the feedback factor is relatively small, a fraction of the output voltage is fed back to the inverted input of the amplifier. This configuration sends the output back to the inverting terminal, creating a negative-feedback condition. If β were fed back to the non-inverting terminal, this small fraction of the output voltage would be added instead of subtracted. In this configuration, there would be positive feedback and the output would eventually saturate.

Closed-Loop Transfer Function

If you analyze the loop in Figure 3, you must assume an output voltage exists. This makes the voltage at A equal to $V_{OUT}(j\omega)$. The signal passes through the feedback system, $\beta(j\omega)$, so that the voltage at B is equal to $\beta(j\omega)V_{OUT}(j\omega)$. The voltage, or input voltage, at C is added to the voltage at B; C is equal to $(V_{IN}(j\omega) - \beta(j\omega)V_{OUT}(j\omega))$. With the signal passing through the gain cell, $A_{OL}(j\omega)$, the voltage at point D is equal to $A_{OL}(j\omega)(V_{IN}(j\omega) - \beta(j\omega)V_{OUT}(j\omega))$. This voltage is equal to the original node, A, or V_{OUT} . The formula that describes this complete closed-loop system is equal to:

$$A = D \text{ or}$$

$$V_{OUT}(j\omega) = A_{OL}(j\omega)(V_{IN}(j\omega) - \beta(j\omega)V_{OUT}(j\omega))$$

By collecting the terms, the manipulated transfer function becomes:

$$V_{OUT}(j\omega)/V_{IN}(j\omega) = A_{OL}(j\omega) / (1 + A_{OL}(j\omega) \beta(j\omega))$$

This formula is essentially equal to the closed-loop gain of the system, or $A_{CL}(j\omega)$.

This is a very important result. If the open-loop gain ($A_{OL}(j\omega)$) of the amplifier is allowed to approach infinity, the response of the feedback factor can easily be evaluated as: $A_{CL}(j\omega) = 1/\beta(j\omega)$

This formula allows an easy determination of the frequency stability of an amplifier's closed-loop system.

Calculation of $1/\beta$

The easiest technique to use when calculating $1/\beta$ is to place a source directly on the non-inverting input of the amplifier and ignore error contributions from the amplifier. You could argue that this calculation does not give the appropriate circuit closed-loop-gain equation for the actual signal, and this is true. But, if you use this calculation, you can determine the level of circuit stability.

The circuits in Figure 4 show how to calculate $1/\beta$.

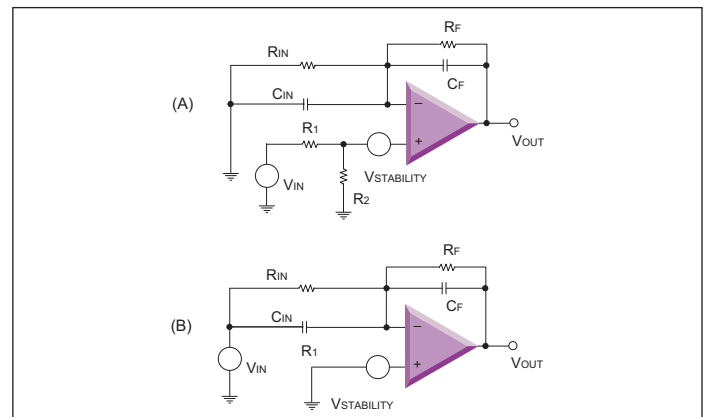


Figure 4: The input signal in circuit a.) at DC is gained by $(R_2/(R_1+R_2))(1+R_f/R_{IN})$. The input signal in circuit b.) has a DC gain of $-R_f/R_{IN}$. Neither of these gain equations match the DC gain of the feedback factor, $1/\beta$.

Operational Amplifiers

In Figure 4A and 4B, $V_{STABILITY}$ is a fictitious voltage source equal to zero volts. It is used for the $1/\beta$ stability analysis. Note that this source is not the actual application-input source.

Assuming that the open-loop gain of the amplifier is infinite, the transfer function of this circuit is equal to:

$$V_{OUT}/V_{STABILITY} = 1/\beta$$

$$1/\beta = 1 + (R_F||C_F)/(R_{IN}||C_1)$$

$$(C_1 = C_{IN} + C_{CM} \text{ for Figure 4A and } C_1 = C_{IN} + C_{CM} + C_{DIFF} \text{ for Figure 4B})$$

or

$$1/\beta(j\omega) = (R_{IN}(j\omega)R_FC_F + 1) + R_F(j\omega)R_{IN}C_1 + 1)/R_{IN}(j\omega)R_FC_F + 1)$$

In the equation above, when ω is equal to zero:

$$1/\beta(j\omega) = 1 + R_F/R_{IN}$$

As ω approaches infinity:

$$1/\beta(j\omega) = 1 + C_1/C_F$$

The transfer function has one zero and one pole. The zero is located at:

$$f_z = 1/(2\pi R_{IN}||R_F(C_1+C_F))$$

$$f_p = 1/(2\pi R_FC_F)$$

The Bode plot of the $1/\beta(j\omega)$ transfer function of the circuit in Figure 4A is shown in Figure 5.




Figure 5: These are the Bode plots of the inverse of the feedback factor ($1/\beta$) for the circuit in Figure 4A, using $V_{STABILITY}$ as the input source.

Once again, in Figure 4B the input source that is used for this analysis is not the same as the input source for the actual application circuit. However, the amplifier stability is determined in the same manner. The closed-loop transfer function, using $V_{STABILITY}$, is equal to:

$$V_{OUT}/V_{STABILITY} = 1/\beta$$

$$1/\beta = 1 + (R_F||C_F)/(R_{IN}||C_1)$$

or

$$1/\beta(j\omega) = (R_{IN}(j\omega)R_FC_F + 1) + R_F(j\omega)R_{IN}C_1 + 1)/R_{IN}(j\omega)R_FC_F + 1)$$

Note that the transfer functions of $1/\beta$ between Figure 4A and 4B are identical.

Determining System Stability

If you know the phase margin, you can determine the stability of the closed-loop amplifier system. In this analysis, the Bode stability-analysis technique is commonly used. With this approach, the magnitude (in dB) and phase response of both the open-loop response of the amplifier and circuit feedback factor are included in a Bode plot.

The system closed-loop gain is equal to the lesser (in magnitude) of the two gains. The phase response of the system is the equal to the open-loop gain phase shift minus the inverted feedback factor's phase shift.

The stability of the system is defined at the frequency where the open-loop gain of the amplifier intercepts the closed-loop gain response. At this point, the theoretical phase shift of the system should be greater than -180 degrees. In practice, the system phase shift should be smaller than -135 degrees. This technique is illustrated in Figures 6 through 9. The cases presented in Figure 6 and 7 represent stable systems. The cases presented in Figures 8 and 9 represent unstable systems.




Figure 6: This closed-loop system is stable with a phase shift of -90 degrees at the intercept of the A_{OL} and $1/\beta$ curves.

In Figure 6, the open-loop gain of the amplifier ($A_{OL}(j\omega)$) starts with a zero dB change in frequency and quickly changes to a -20 dB/decade slope. At the frequency where the first pole occurs, the phase shift is -45 degrees. At that frequency, one decade above the first pole, the phase shift is approximately -90 degrees. As the gain slope progresses with frequency, a second pole is introduced, causing the open-loop-gain response to change -40 dB/decade. Once again, this is accompanied with a phase change. The third incident that occurs in this response is where a zero is introduced and the open-loop gain response returns back to a -20 dB/decade slope.

The $1/\beta$ curve in this same graph starts with a zero dB change with frequency. This curve remains flat with increased frequency until the very end of the curve, where a pole occurs and the curve starts to attenuate -20 dB/decade.

Operational Amplifiers

The point of interest in Figure 6 is where the $A_{OL}(j\omega)$ curve intersects the $1/\beta$ curve. The rate 20 dB/decade of closure between the two curves suggests the phase margin of the system and in turn predicts the stability. In this situation, the amplifier is contributing a -90 degree phase shift and the feedback factor is contributing a zero-degree phase shift. The stability of the system is determined at this intersection point. The system phase shift is calculated by subtracting the $1/\beta(j\omega)$ phase shift from the $A_{OL}(j\omega)$ phase shift. In this case, the system phase shift is -90 degrees. Theoretically, a system is stable if the phase shift is between zero and -180 degrees. In practice, you should design to a phase shift of -135 degrees or smaller.




Figure 7: This system is marginally stable with a -135 degree phase shift at the intersection of the two gain curves.

In the case presented in Figure 7, the point of intersection between the $A_{OL}(j\omega)$ curve and the $1/\beta(j\omega)$ curve suggests a marginally stable system. At that point the $A_{OL}(j\omega)$ curve is changing -20 dB/decade. The $1/\beta(j\omega)$ curve is changing from a +20 dB/decade to a 0 dB/decade slope. The phase shift of the $A_{OL}(j\omega)$ curve is -90 degrees. The phase shift of the $1/\beta(j\omega)$ curve is +45 degrees. The system phase shift is equal to -135 degrees.

Although this system appears to be stable, i.e., the phase shift is between zero and -180 degrees, the circuit implementation is not as clean as calculations or simulations would imply. Parasitic capacitance and inductance on the board can contribute additional phase errors. Consequently, this system is “marginally stable” with this magnitude of phase shift. This closed-loop circuit has a significant overshoot and ringing with a step response.

In Figure 8, the $A_{OL}(j\omega)$ is changing at a rate of -20 dB/decade. The $1/\beta(j\omega)$ is changing at a rate of +20 dB/decade. The rate of closure of these two curves is 40 dB/decade and the system phase shift is -168 degrees. The stability of this system is very questionable.

In Figure 9, $A_{OL}(j\omega)$ is changing at a rate of -40 dB/decade. The $1/\beta(j\omega)$ is changing at a rate of 0 dB/decade. The rate of closure in these two curves is 40 dB/decade, indicating a phase shift of -170 degrees. The stability of this system is also questionable.




Figure 8: In a practical circuit implementation, given layout parasitics, this system is unstable.




Figure 9: In a practical circuit implementation, given layout parasitics, this system is also unstable.

Conclusion

At the beginning of this article you were asked, “What is operational-amplifier circuit stability and how do you know when you are on the “hairy edge?” There are many definitions of stability in analog, such as, unchanging over temperature, unchanging from lot to lot, noisy signals, etc. But, an analog circuit becomes critically unstable when output unintentionally oscillates without excitation. This kind of stability problem stops the progress of circuit design until you can track it down.

You can only evaluate this kind of stability in the frequency domain. A quick paper-and-pencil examination of your circuit readily provides insight into your oscillation problem. The relationship between the open-loop gain of your amplifier and the feedback system over frequency quickly identify the source of the problem. If you use gain and phase Bode plots, you can estimate where these problems reside. If you keep the closed-loop phase shift below -135 degrees your circuit oscillations do not occur and ringing will be minimized. If you do this work up front with amplifiers, you can avoid those dreadful designs that kick into an unwanted song, a.k.a “The Amplifier Circuit Blues”.

Amplifiers And The SPICE Of Life

By Bonnie C. Baker, Microchip Technology Inc.

What is operational amplifier circuit stability and how do you know when you are on the hairy edge? Typically, op amps are used with a feedback network in order to reduce the variability of the open loop gain response from part to part. With this technique, circuit stability is provided. But it is possible to design an amplifier circuit that does quite the opposite. You can design an amplifier circuit that is extremely unstable to the point of oscillation. In a closed loop amplifier system, stability can be determined if the phase margin of the system is known. In this analysis, the Bode stability analysis technique is commonly used. With this technique, the magnitude (in dB) and phase response of both the open loop response of the amplifier and circuit feedback factor are included in a Bode plot.

They say a computer-based simulation of your analog circuit is important. This is because the use of your preferred computer SPICE (Simulation Program with Integrated Circuit Emphasis) program can reduce initial errors and development time. If you use your SPICE simulator correctly, you can drum out circuit errors and nuances before you go to your breadboard. In this manner, you will verify your design before you spend the time to solder your circuit. SPICE helps troubleshoot on the bench; it is a great place to try out different hypotheses. It is great at “what if” scenarios (i.e., exploratory design).

You can view the results from these software tools on a PC with user-friendly GUI suites. This tool will fundamentally provide DC operating (quiescent) points, small signal (AC) gain, time domain behavior and DC sweeps. At a more sophisticated level, it will help you analyze harmonic distortion, noise power, gain sensitivity and perform pole-zero searches. This list is not complete, but generally, SPICE software manufacturers have many of these fundamental features available for the user. By finessing the Monte Carlo and worst-case-analysis tools in SPICE, you can predict the yields of your final product. If you use your breadboard for this type of investigation, it could be very expensive and time consuming. All of these things will speed up your application circuit time-to-market.

But, beware. You can effectively evaluate analog products if your SPICE models or macromodels are accurate enough for your application. The key words here are “accurate enough”. Such models, or macromodels, should reflect the actual performance of the component, without carrying the burden of too many circuit details. Too many details can lead to convergence problems and extremely long simulation times. Not enough details can hide some of the intricacies of your circuit’s performance. Worse yet, your simulation, whether you use complete models or just macromodels, may give you a misrepresentation of what your circuit will really do. Remember that a SPICE simulation is simply a pile of mathematical equations that, hopefully, represent what your circuit will do. It is in essence a computer product that produces imaginary results.

So you might ask, “why bother?” Are SPICE simulations worth the time and effort? A pop quiz will help you clarify this question. The circuit in Figure 1 shows a fundamental, basic circuit. Is this circuit stable or does it oscillate? Would the output of the amplifier have an unacceptable ring? I would think that you would quickly look at this and say, “That is a silly question. Of course it is stable!” But then again, if you are always looking for the trick question you may be suspicious. So what is the answer?




Figure 1: A variety of applications across industry has this simple sub-circuit embedded in the system. This circuit simply takes an analog input signal and gains that signal to the output of the amplifier. For instance, an input signal of +0.5V to +1 VDC would become a +1V to +2 VDC signal at V_{OUT}. The question is, would this DC signal oscillate? Or, would a 50 kHz sinusoidal signal oscillate or ring? The bandwidth of this amplifier is 2.8 MHz.

This simple amplifier circuit uses an amplifier in a gain of +2 V/V. The amplifier has an 100 kΩ resistor connected to its inverting-input to ground and 100 kΩ resistor in the feedback loop. It would be easy to assume that this circuit is stable. However, tedious calculations will verify that this amplifier circuit will ring. This is due to the parasitic capacitances around the resistors and the high differential/common-mode capacitance of the amplifier’s input stage. For this particular amplifier, the input common-mode capacitance is 6 pF and the differential-mode capacitance is 3 pF. These capacitances interact with the feedback resistor causing a semi-unstable condition. If you bench-test this circuit you will immediately see this condition on the oscilloscope. Parasitics on the breadboard will aggravate this instability.




Figure 2: By enhancing the circuit in Figure 1 with the parasitic capacitances of the resistors and amplifier, a simple of a circuit is not so simple. In the DC domain these capacitors will operated as open circuits. In the AC domain, the capacitors will affect the perfect square wave from the input to output. The perfect square wave will have quite a ring at the V_{OUT} node.

Operational Amplifiers

Enhancing the circuit in Figure 1 with the parasitic capacitances of the resistors and amplifier, things aren't as simple. The 100 k Ω resistor in the feedback loop will have a parallel 0.5 pF (approximately) and the parasitic to ground could be as high as 2 pF or 3 pF. In the dc domain these capacitors will operate as open circuits. In the ac domain the capacitors will affect the perfect square wave from input to output and there will be quite a ring at the Vout node. If you use the amplifier's SPICE macromodel, with input impedances in the model and board parasitics, you will see this problem immediately in your simulation. If you bread board the circuit, you most certainly will see this ringing.

Changing the values of the two resistors in this circuit solves this problem. Hand calculations will help you find the correct values. A SPICE simulation will facilitate the process. This is a little easier than swapping out resistors on the breadboard until you find the right values. In SPICE, you can also look at the response of the amplifier using various resistors. This will help you find the "corner" of this oscillation. If you go back and change both values to 10 k Ω you will have great success in SPICE and on the bench. Figure 3. shows the simulation results.

You do need to understand where the problem came from before you continue with your circuit design. but the simulation caught a significant stability problem which was an easy one to miss by just an inspection of the schematic.

The nay say-ers in industry will tell you your computer based simulation tools will not work and using them will be a waste of time. These people are a bit misguided and in my opinion, they have a superficial view of what this tool can really do. Sure, SPICE tools can lead you astray. But like any tool, it is only as good as the user. Any insight that you gain from your simulations emerges if you understand and use you SPICE tools properly. Better yet, SPICE simulations will point out problems that you had never anticipated. In most cases, they use double precision calculations. This makes it easier to detect low level problems that are impossible to find on the bench. SPICE simulations of analog circuits is just one more way to avoid the design of a "singing" circuit when you actually want stability.

This article is excerpted from Bonnie Baker's book, "A BAKER'S DOZEN: Real Analog Solutions for Digital Designers," published in Spring 2005 by Elsevier.

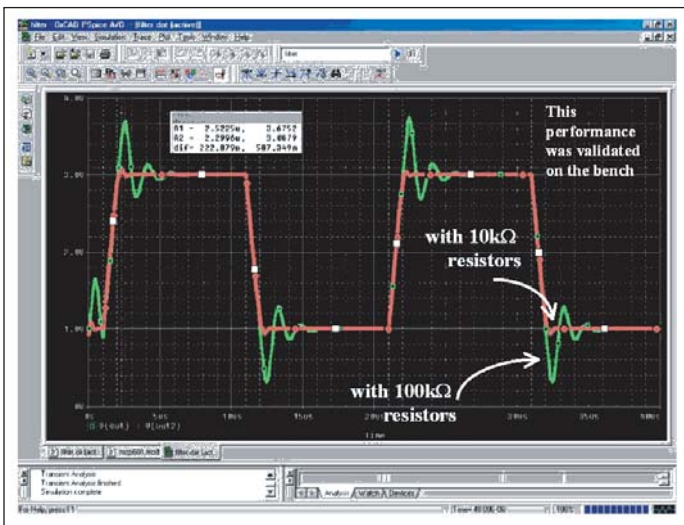


Figure 3: You can quickly verify that this simple circuit will ring using a SPICE simulation. If you need to double-check this with a breadboard circuit that is also a good idea, however, reducing the 100 k Ω resistors down to 10 k Ω resistors solves the problem. You do need to understand where the problem came from before you continue with your circuit design. But this simulation caught a significant stability problem. This ringing problem was an easy one to miss by inspection of the schematic.

The Number Of Bits vs LSB Errors

By Bonnie C. Baker, Microchip Technology Inc.

What do the Least Significant Bits (LSB) specifications mean when you are looking at Analog-to-Digital converters (ADC)? A fellow engineer told me that a 12-bit converter, from X manufacturer, had just seven usable bits. So, essentially the 12-bit converter was only a 7-bit converter. He based this conclusion on the device's offset and gain specifications. The maximum specifications were:

offset error = $\pm 3 \text{ LSB}$,

gain error = $\pm 5 \text{ LSB}$,

At first glance, I thought he was right. From the list above, the worst specification is gain error ($\pm 5 \text{ LSB}$). Applying simple mathematics, 12 bits minus 5 bits of resolution is equal to 7 bits, right? Why would an ADC manufacturer introduce such a device? The gain-error specification motivates me to purchase a lower-cost, 8-bit converter, but that doesn't seem right. Well, as it turns out, it wasn't right.

Let's start out by looking at the definition of LSB. Think of a serial 12-bit converter, it produces a string of 12 ones or zeros. Typically, the converter's first transmitted digital bit is the Most Significant Bit (MSB) (or $\text{LSB} + 11$). Some converters transmit the LSB first. We will assume that the MSB is first in this discussion (shown in Figure 1). The second bit is MSB-1 (or $\text{LSB}+10$), the third bit is MSB-2 (or $\text{LSB}+9$), etc. At the end of this string of bits, the converter finally transmits as MSB-11 (or LSB).

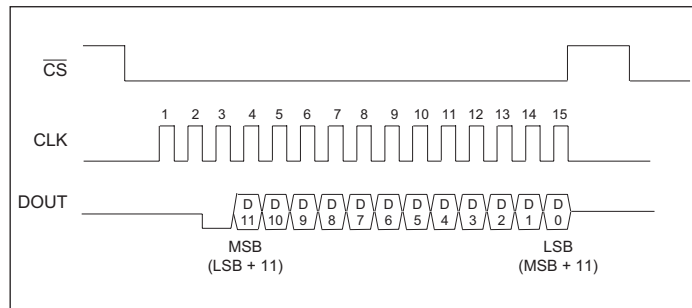


Figure 1: The data from this serial ADC clocks MSB out first and LSB last.

The terminology, LSB, is very specific. It describes the last position in the digital stream. It also represents a fraction of the full-scale input range. For a 12-bit converter, the LSB value is equivalent to the analog full-scale input range, divided by 2^{12} or 4,096. If I put this in terms of real numbers, I have an LSB size of 1 mV with a 12-bit converter that has full-scale input range of 4.096V. However, the most instructive definition of LSB is that it can represent one code out of the 4096 codes possible.

Going back to the specifications and translating them into a 12-bit converter that has an input full-scale range of 4.096V:

offset error = $\pm 3 \text{ LSB} = \pm 3 \text{ mV}$,

gain error = $\pm 5 \text{ LSB} = \pm 5 \text{ mV}$,

These specifications actually claim that the converter can have (worst case) an 8 mV (or 8 code) error introduced through the conversion process. This is not to say that the error occurs at the LSB, LSB-1, LSB-2, LSB-3, LSB-4, LSB-5, LSB-6 and LSB-7 positions in the output-bit stream of the converter. The errors can be up to eight times one LSB, or 8 mV. Precisely stated, the transfer function of the converter could have up to eight codes missing out of 4,096 codes. These codes will be missing at the lower or upper range of the codes. For instance, a converter with an error of +8 LSB ((+3 LSB offset error) + (+5 LSB gain error)) will produce possible output codes of zero to 4,088. The lost codes are from 4088 up to 4,095. This is a small, incremental error of 0.2% at full-scale. In contrast, a converter with an error of -3 LSB ((-3 LSB offset error) - (-5 LSB gain error)) will produce codes from three up to 4,095. The gain error in this situation produces an accuracy problem, not a loss of codes. The lost codes are 0, 1 and 2. Both of these examples illustrate the worst possible scenario. Typically, the offset errors and gain errors do not track this closely in actual converters.

The real-life performance enhancements, due to incremental improvements in an ADC's offset or gain specifications, are negligible to non-existent. To some designers this seems like a bold assumption, if precision is one of the design objectives. It is easy to implement digital calibration algorithms with your firmware. However, more importantly, the front-end amplification/signal conditioning section of the circuit typically produces higher errors than the converter itself.

This discussion puts a new light on the conclusions reached at the beginning of this article. In fact, the 12-bit converter, as specified above, has an accuracy of approximately 11.997 bits. The good news is that a microprocessor or microcontroller can remove this offset and gain error with a simple calibration algorithm.

Analog-to-Digital Converters

Is Your Converter Accurate?

By Bonnie C. Baker, Microchip Technology Inc.

What does analog-to-digital converter (ADC) accuracy really mean? You might say that accuracy means the output of the ADC code represents the actual analog-input voltage (minus the quantization error). This makes sense, but are you seeing a precise determination of the analog input? Does accuracy also mean that the ADC conversion results are repeatable? Is the converter output code repeatable from transition to transition, with everything remaining unchanged in the circuit during successive conversions?

In your mind, an ADC code-to-code transition may be sharp, occurring at a unique input voltage. Actually, the transition regions in the ADC-transfer function may be wide. In fact, these regions may span across several digital-output codes. In Figure 1, a transition point occurs when the digital output switches from one code to the next with respect to a specific analog input voltage. But, because of ADC internal noise, the transition point is typically not a single threshold, but rather a small region of uncertainty. Consequently, you need to define the transition point as the statistical average of many conversions. Stated differently, it is at the voltage input where the uncertainty of multiple conversions averages 50 percent of the time to one digital code, and 50 percent of the time to an adjoining digital code. Upon closer inspection, the conversions you collect appear to be noisy, with a gaussian probability curve.




Figure 1: This non-ideal transfer function of an 8-bit ADC illustrates the transition noise of every code. In this graph, FS means full-scale.

One experiment that you can quickly run with your converter is to ground the input of a good 16-bit ADC. I am assuming, of course, that you are using good layout techniques, bypass capacitors, etc. in your circuit. Now, collect 1,024 samples at the converter's specified conversion rate. You will find that you have multiple codes in your output data. What you will witness is the transition noise of the converter. Some manufacturers will tell you what the rms transition noise is for their ADCs. Multiply the rms transition noise specification by 6.6 to obtain a peak-to-peak value.

Let's take this discussion a step further. The offset, gain, differential-non-linearity (DNL) and integral-non-linearity (INL) are the accuracy specifications for ADCs. Some manufacturers also call these the DC specifications. This is because these device tests use a DC-input voltage for the conversions. But, these specifications do not tell you how repeatable the results are from conversion to conversion. They only tell you that, on the average, these errors will be no more or less than the minimum and maximum in your ADC manufacturer's data sheet. In order to describe the accuracy of your converter precisely, you need to combine the AC specifications with the DC specifications.

Within the AC specifications, you will find three types. However, one in particular is interesting for this discussion. This AC specification is called the (Signal-to-Noise + Distortion) ratio, otherwise known as SINAD. The counterpart to this specification is "effective number of bits" (ENOB). $ENOB = (SINAD - 1.76dB) / 6.02$. This specification, combined with the DC specifications, will give you a stronger feel for how accurate your converter really is.

References

- "Anticipate the Accuracy of Your Converter", Bonnie C. Baker, EDN, March 18, 2004.
- "Using the Analog-to-Digital (A/D) Converter", Mitra, D'Souza, Cooper, AN546, Microchip Technology Inc.
- "Analog-Digital Conversion Handbook", David Sheingold, Prentice-Hall, 1986.

This article is excerpted from Bonnie Baker's book, "A BAKER'S DOZEN: Real World Solutions to Real World Analog Design Problems," published Spring 2005 by Elsevier.

Sometimes Noise Can Be Good

By Bonnie C. Baker, Microchip Technology Inc.

When I started to write this article someone saw the title and asked, “You don’t have any kids, do you?” Well, I did have the charge of children at one time in my life, but I think people forget two things through their lives: extreme pleasure and extreme pain. That’s why we keep going back for more. This column is not about going back to painful experiences such as a noisy circuit. Rather, it is about the pleasure of tackling those difficult analog noise problems in the digital domain.

We have all sought the perfect conversion in our mixed-signal circuits where the converter produces a repeatable, accurate digital result every time. We use noise-reduction techniques such as selecting low-noise devices, a careful layout and analog filtering to remove undesirable signals. But, another way to approach noisy analog-to-digital conversion problems is to “design” noise into your signal instead of out. For instance, you can get 12-bit accuracy from a 12-bit converter if you are diligent about applying low-noise strategies to your circuit. As an alternative, you can allow a degree of white noise into the circuit and follow the conversion with a processor or controller digital filter. In this scenario, your circuit is capable of producing 14-, 15- or even 16+-bit accuracy. If there is noise in your circuit, you can achieve better resolution at the output of a digital filter by using oversampling techniques.

For instance, if you use a simple rolling-average digital filter, you can calculate the number of bits (N) that you will add to your conversion resolution with this formula: $\#oversampled\ data = 2^{2N}$. If you want to increase your resolution from 11 bits to 14 bits, you need to accumulate and average 64 samples. Time is the primary tradeoff for this increase in resolution. The rolling-average digital-filter algorithm accumulates several samples in order to calculate the final result. The accumulation of these samples takes time. Finite Impulse Response (FIR) and Infinite Impulse Response (IIR) digital filters are also well suited for this task.

If you have the time, this sounds like a simple solution. However, there is one more issue to consider before you embrace this combination of analog with digital systems in your circuit. You must know the complexion of your ADC digital output over time. A histogram plot is an appropriate tool to use when examining your digital code over time. The histogram plot displays the number of occurrences of each code. For example, the histogram plot in Figure 1 shows 1024 repetitive data samples from a 12-bit ADC (sample rate = 20 ksp/s).

If you want to successfully increase the resolution of your converter, you need to ensure that the noise from the ADC is gaussian in nature. In a histogram plot, gaussian noise looks similar to a statistically normal distribution around a center code. The data in Figure 1 does not follow the shape of a normal distribution. The data in Figure 1 appears to have a bi-modal response. In addition, the output mean of this system should be 2236 instead of 2297. A digital filter will not “fix” this data. The noise in this system originates in an LED array. Poor layout and high currents through the array make the noise on the board intolerable.

If you use a digital filter at the output of your ADC, you are not relieved of the responsibility of knowing what kind of data you are producing. Digital filtering will improve the resolution of your analog-to-digital conversion, but only if you are confident that the noise response of your data is gaussian in nature.




Figure 1: This histogram plot shows data from an ADC that is bimodal. It is difficult, if not impossible, to apply digital filtering to this data to recover a digital code that actually represents the input analog voltage accurately.

References

“Mixed-Signal and DSP Design Techniques”, Walt Kester, Elsevier, 2003.

Analog-to-Digital Converters

Using The Basics For ADC Range Control

By Bonnie C. Baker, Microchip Technology Inc.

You can use a high-resolution ADC to simplify or eliminate most of the sensitive analog front-end circuitry in your system. I know that this sounds too good to be true, but consider the pressure-sensing example of a single-supply signal path that measures pressure from the real world to a controller. The electronic portion of this signal path starts at the pressure sensor. The signal then travels through an instrumentation amplifier (the gain stage with a voltage reference) into a 5th order analog low-pass filter (i.e., the two op amp noise reduction stage), digitized by a 10- or 12-bit ADC and finally into the microcontroller. In the microcontroller or processor, you finally implement code that calibrates the data and reduces error. Wow, can you say all of that in one breath? At minimum, this signal path requires seven active devices.

This is a long-winded approach, especially when you consider a high-resolution converter, such as a 24-bit sigma-delta converter, as an alternative. With the sigma-delta converter, the only external circuitry before the converter is the pressure sensor and a few 1st order, R/C filters. This is a great strategy because you can reduce noise and layout errors with the sigma-delta converter. You can also calibrate system-offset and gain errors while getting your 12-bit range back in the controller or processor.

Normally, you would adjust the 24-bit converter's offset error at the code transition between 000h and 001h. Since converters have some degree of transition noise, you need to sample this transition area several times to verify the transition voltage. You can then quickly remove the offset error of the converter using the following formula:

$$\text{Offset error} = (V[0:1] - 0.5 (V_{\text{LSB}})) / V_{\text{LSB}}$$

$$V_{\text{LSB}} = V_{\text{REF}} / 2^n = \text{ideal LSB voltage size}$$

$$V[0:1] = \text{analog voltage of first transition}$$

$$V_{\text{REF}} = \text{full-scale voltage}$$

$$n = \text{number of converter bits}$$

This formula is a good start. We can translate it into a usable form to operate at any point in the ADC transfer function (also see Figure 1).

$$\text{Offset error (with level shift)} = (V[x:(x+1)] - 0.5 V_{\text{LSB}}) / V_{\text{LSB}}$$

"x" is the code that is produced by your tested or selected offset value.

With this new offset-error formula, you can inject a level shift into the output data with the controller. From this point, you can calculate the gain error for the region of interest. The formula that calculates the gain error across the full-scale range of an ADC is:

$$\text{Gain error} = (V_{\text{REF}} - 2 V_{\text{LSB}} - V[(2^n - 2):(2^n - 1)] - V[0:1]) / V_{\text{LSB}}$$

If you translate this gain error formula to match the range of interest, it becomes:

$$\text{Gain error} = (V_{\text{REF}} - 2 V_{\text{LSB}} - V[(2^N - 2):(2^N - 1)] - V[0:1]) / V_{\text{LSB}}$$

Where capital "N" is equal to the number of bits that you are going to use in your system.

You can apply the offset and gain equations above to the output data of any ADC, regardless of the converter's resolution. High-resolution ADCs can eliminate amplifier gain stages, high-order filter stages and level-shift circuitry by taking advantage of a portion of the million-plus possible output codes. With the offset and gain formulas above, you can use the controller or processor computation power to center on the portion of the conversion that is of interest for your application.




Figure 1: You can isolate the usable output-code range of an ADC by using modified offset and gain error equations.

Turning Nyquist Upside Down By Undersampling

By Bonnie C. Baker, Microchip Technology Inc.

Bravo to Harry Nyquist and Claude Shannon. According to these two gentlemen, in the 1920s, the well-known Nyquist theorem was created stating, “when sampling a signal at discrete intervals, the sampling must be *greater than twice* the highest frequency of the input signal.” This is done so that you can reconstruct the original signal perfectly from the sampled version.

Their theorem has held up well through the ages. In my discussions with engineers, I use the Nyquist theorem to explain the accuracy of sampling systems where the bandwidth of the signal of interest is less than twice the sampling frequency of the converter. The sampling systems I describe use a low-pass, anti-aliasing filter before the ADC. This is usually an engineer’s initial exposure to the Nyquist theorem, where signals with frequencies greater than $\frac{1}{2}$ of the sampling rate of a converter will come back to haunt you. The experts fondly call this a “fold-back” phenomena⁽¹⁾. If the sample rate is less than twice the maximum-input frequency, the digitizing system will produce a mixture of in-band and out-of-band data. Once this fold-back of signal information has occurred, there is no going back in terms of retrieving the original signal that is below $\frac{1}{2}$ of the sampling frequency. So, my advice for this type of system is to always place a low-pass, anti-aliasing filter before your ADC.

That is all well and good, but let’s try to turn the Nyquist theorem upside down. We can use this theorem in an alternate manner by intentionally forcing a system configuration that aliases or folds back higher-frequency signals that occur above the sampling rate of the converter. This is known as undersampling. Some synonyms for undersampling are bandpass sampling or super-Nyquist sampling. Applications, such as wireless communication receivers, radar instrumentation, infrared instrumentation or video, are well suited for this use of the Nyquist theorem.

In these systems, the bandwidth of the signal of interest (Δf_{SIG}) is centered at a higher frequency than the sampling frequency (f_{SAMPLE}) of the converter. Δf_{SIG} is also riding on a high-frequency carrier signal (f_{CAR}). Δf_{SIG} is limited by an analog bandpass filter that acts like an anti-aliasing filter in this system. It is not unusual to implement a simple second-order filter (one zero and one pole) for this purpose. The order and response of this filter is user defined. If you design in a higher order filter, the bandwidth of Δf_{SIG} is smaller.

Two formulas will help you determine the sampling frequency of your system. The first equation is $f_{SAMPLE} > 2(\Delta f_{SIG})$. This formula complements the Nyquist theorem directly. The second formula is $f_{SAMPLE} = 4 f_{CAR} / (2 * Z - 1)$, where Z is a whole number rounded down. You will use this second formula twice as you zero in on the actual sampling frequency.

With the first formula above, the sampling frequency should be equal to twice Δf_{SIG} . Then, by using this calculated value for the sample frequency and a predetermined carrier frequency, you can calculate the quantity of Z in the second formula. The value of Z is usually not a whole number and should be rounded down. With this new value of Z, you should use the second formula to recalculate a value for f_{SAMPLE} .

An example may clarify any questions that you have. For instance, if a system has a signal that has a 3.5 MHz bandwidth (Δf_{SIG}) that is centered at a 70 MHz carrier frequency (f_{CAR}), initially you can calculate the sampling frequency as 7 Msps (f_{SAMPLE} , formula #1). With this number for f_{SAMPLE} , the calculated value for Z is equal to 20.5 (formula #2). Rounding down, Z is equal to 20 making the actual sampling frequency (f_{SAMPLE}) equal to 7.18 MHz.

Beyond the selection of your sampling clock, there are several important issues to think about with your undersampling application. You should select an ADC with an input stage that can accept signals with frequencies above the converter’s sampling rate. An undersampling converter’s product data sheet will specify this. Jitter and phase noise of the converter’s sample clock (f_{SAMPLE}) can degrade the system performance. You may also require a high-quality crystal oscillator.

This column gives you a short tour of undersampling theory. If you need more information, refer to the references below.

References

“Filtering? Before or after?”, Bonnie C. Baker, EDN, February 20, 2003.

“The Data Conversion Handbook”, Walt Kester, Elsevier, ISBN0-506-7841.

“Putting Undersampling to Work”, Pentek, Inc.
www.pentek.com/applications

Miscellaneous Articles

Arming Yourself With nanoWatt Technology Techniques

By Bonnie C. Baker, Microchip Technology Inc.

We have come to expect and require more from our battery-powered equipment. My first Personal Digital Assistant (PDA) would only retain its battery power for a day. If the calendar alarm was enabled, the battery power would quickly go to nothing. Today, my PDA holds its battery charge for an entire week under the same conditions. Both PDAs had the same battery chemistry, Lithium Ion (Li Ion), with the same power density. So what has changed? Simple. The hardware was improved and the power-management techniques were refined. Battery improvements were secondary.

PDAs are a great example of where system sophistication is improving rapidly, but as fate would have it, the battery-powered equipment requirements are increasing while battery chemistries remain the same. You can accomplish this increase of functionality only if the firmware/software engineers understand the tools available in the microcontroller-based processing unit and if the hardware designers understand the efficiency of the solutions available on the market today.

Making Analog and Digital Play Together

You can accomplish an improvement in power consumption, and consequently an increase in functionality, if you understand the hardware options and the tools available in microcontrollers. One dimension of power conservation is controlling the magnitude of the power-supply voltage in your application. You are probably interfacing to the real world at some point during your code operation. If you are, you will have analog content in your circuit. Analog power-supply requirements are higher than the requirements for digital. Don't forget that analog-noise margins are much smaller than digital-noise margins. The analog noise floor does not reduce with lower power-supply voltages. It stays the same over power-supply-voltage changes. For instance, a 12-bit Analog-to-Digital Converter (ADC) can produce good, solid conversions with a 5V supply. However, that same 12-bit ADC will produce a smaller number of noise free-bits when you use a 2V supply. This is because the least significant bites (LSb) size has become smaller, but the magnitude of the noise is consistent. The solution to this problem is to use higher supply voltages when running analog and lower voltages during the digital-only operation.

The diagram in Figure 1 shows a simple, microcontroller-based, battery-operated system using the PIC18F1320 Flash microcontroller from Microchip Technology Inc. The PIC18F1320 has features, such as a variety of idle modes and a two-clock start-up capability, that can enhance your low-power strategy.

On the hardware side, the industry is continuing to develop classes of external peripherals as well as the internal microcontroller peripherals, with lower power performance in mind. In terms of the external peripherals to the microcontroller, you can achieve lower power by reducing power-supply-voltage requirements to the chips and optimizing topologies for the lower-power jobs. This simple example (Figure 1) has low-power operational amplifiers, an ADC and a regulated, adjustable charge pump.




Figure 1: The lower-power, external peripherals to the microcontroller only provide half of this battery-powered system implementation. The programmable capabilities of the microcontroller allow control of the power-supply voltage in the system (MCP1252-ADJ) and control of the microcontroller system clock and sleep modes.

The design of the operational amplifiers in Figure 1 uses CMOS. This type of operational amplifier is continuing to push minimum power-supply voltage requirements down. The MCP6041 from Microchip, in Figure 1, is a 14 kHz, 600 nA amplifier and requires a supply voltage as low as 1.4V and up to 5.5V. The combination of reduced supply-voltage and lower quiescent current provides a good solution for power management concerns in battery-operated equipment.

With internal or external integrated ADCs, the amount of power dissipated is more dependent on the converter topology than on IC design innovation. For instance, the ratio of conversion time-to-current consumption in the SAR (Successive Approximation Register) converter is considerably lower than in the delta-sigma converter. You will probably use the SAR converter in battery-powered applications, unless you need higher resolution and accuracy.

The power supply in the circuit in Figure 1 is adjustable. A higher voltage of 5V is best suited for analog circuitry and a lower voltage of 2V is best suited for digital activities. The adjustable power-converter in Figure 1 has high efficiency with low-output currents and Li-Ion battery input voltages (4.2V down to 2.8V). For these reasons, this circuit uses a regulated, adjustable charge pump, DC/DC converter (MCP1252-ADJ).

Controlling the power-supply voltage for various operations is only half of the story. If you really have a lower power "state of mind", you will want to power down some parts of the microcontroller while letting other sections continue, to operate. As an example, you can independently run an A/D or D/A conversion or the USART communication interface, from the microcontroller. These device functions may only need power locally.

Optimizing the external-peripheral power trade offs is also important. In addition, you will find real power savings when using the external and internal peripherals in concert with the microcontrollers programming capability. For instance, the microcontroller controls the power-supply voltage by switching a new configuration into the resistive feedback system of the MCP1252-ADJ regulated, adjustable-output charge pump. The charge pump generates a higher output voltage to insure that the analog circuitry performs at its optimum level. Digital events from the microcontroller can tolerate a lower power-supply voltage. For instance, the power-supply specifications of the PIC18F1320 are from 2V to 5.5V. You can calculate the power savings for this type of change as a direct ratio of the two voltages from the charge pump. Power savings are further improved if the external peripherals are powered down with the lower-power supply voltage using the I/O ports.

Controlling Your Clocks

One issue that is often overlooked when designers are trying to reduce the overall power consumption of an embedded-system circuit is the management of the clock when the microcontroller comes out of its sleep mode.

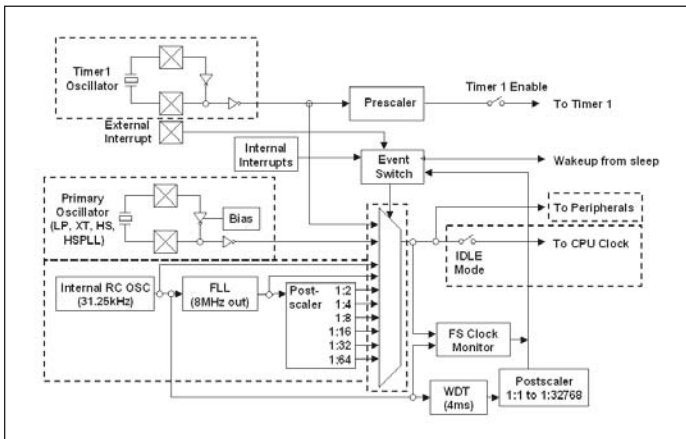


Figure 2: This diagram illustrates an example of three oscillator sources for a microcontroller. Externally, there are two clock connections (Timer 1 and Primary Oscillator). Internally, there is an oscillator.

A microcontroller can have a variety of clock sources (Figure 2). The most obvious clock source is an external one. In this instance, you would connect a crystal oscillator, ceramic resonator, an internal controller clock, or a clock generator to the appropriate device pin. Beyond these elements that generate the clock signal the microcontroller can have a postscaler, prescaler Frequency-lock-loop (FLL). The postscaler and prescaler divide the input-clock frequency down. You would use the FLL to multiply the input-clock frequency.

In a Real Time Operating System (RTOS), it is critical to have a clock management strategy when the system wakes up for short periods of time and sleeps for a long period after the wake-up. If the wake-up time is typically <1 second and you are using a crystal oscillator or ceramic resonator, you may find that there will be a delay between pulling out of the sleep mode and beginning to execute code. The microcontroller will not execute code during this delay or start-up time. However, the application circuit will be consuming power.

For example, Figure 3 shows the typical start-up time for a 4 MHz crystal oscillator. In Figure 3, this time is approximately 450 msec. If this crystal oscillator were the only clock connected to the microcontroller and one second was allocated for the code execution, the actual execution time of the code would be 45% longer than expected. During the clock start-up time your circuit is consuming power but not executing code.




Figure 3: The clock connected to the microcontroller is a 4 MHz LP crystal. In this oscilloscope plot, Channel 1 is a microcontroller output pin. The programming of the microcontroller toggles this output pin every 12 instruction cycles. Channel 2 shows the power-supply to the microcontroller, V_{DD}. The external clock to the microcontroller becomes active approximately 450 msec after power-supply start-up.

With this type of application, it might be advisable to use an internal clock to execute code. The internal clock will start up nearly instantaneously. It is not unusual for a 4 MHz internal clock to start in a few microseconds. Figure 4 illustrates the start-up time of an internal clock.




Figure 4: The start-up time of the internal clock is approximately 1.15 microsec. This is considerably faster than the start-up time of the oscillator (in Figure 3).

Miscellaneous Articles

There is a ~50,000X improvement from the 4 MHz crystal oscillator. From this data, one might summarize the appropriate clock for this type of application is an internal clock. The power consumption of an internal clock is nearly equivalent to the power consumption of a crystal oscillator. This strategy works, as long as your microcontroller is not required to run time-critical operations, such as USART communications or timing a precision pulse.

A third clock source that you may be evaluating is the resonator. Figure 5 illustrates the start-up time of a resonator.




Figure 5: The start-up time of a 4 MHz resonator is faster than a 4 MHz crystal and slower than a 4 MHz internal clock.

There is a clocking system that you can use that is better than any of these three clocking options. The best of all worlds is to quickly determine if your circuit needs a precision clock. If the microcontroller needs a precision clock, the oscillator or resonator is turned on. If not, the microcontroller will shut down. This determination is made quickly, after the microcontroller leaves its sleep mode. If you combine the internal clock with an external resonator or crystal oscillator, this type of decision can quickly be made. You can find significant improvements in power consumption when you use two clock sources (instead of one).

This technique is called the “Two-Clock Start-Up Strategy”. In this hardware/firmware configuration, the microcontroller uses two clocks. Both clocks are off during the sleep mode of the application. At the time of wake-up, the internal clock is turned on to quickly determine if the crystal oscillator is required. If it is required, the clock continues to execute code until the crystal oscillator is up and running. At this time, the microcontroller switches over to the crystal oscillator and turns off the internal clock.

Working the Digital Angle with Sleep Modes

The central focus of a successful low-power design is a microcontroller that has a variety of sleep modes and clock modes. You can conserve system power with the idle modes and sleep modes of the microcontroller. The idle modes of the microcontroller power down the CPU while allowing functions such as the 10-bit ADC to continue to operate. The sleep mode implements a complete shutdown of the microcontroller.

When the clock of your microcontroller switches states the various logic gates in your microcontroller pull current from the power source. When looking at the current consumption in your microcontroller, the first stop is to look at the clock power consumption. If you examine the types of clocks available to you, the internal oscillator will run with less power than the frequency equivalent crystals, oscillators or resonators.

Some microcontrollers have three fundamental modes of operation. The first is the full-bore run mode, where everything is up and running. An intermediate mode is the idle or wait mode, where the peripherals are usually running but not the microcontroller. The third and most important mode for lower-power, battery operation is the sleep or stop mode. In this sleep mode, the device stops consuming power completely. The sleep mode generally disables the system’s clocks, but power conservation is more effective if you also disable the external clock sources.

Here are some additional suggestions to complete your low-power strategy. Drive any unused I/O pins into a high or low state. Use the internal oscillators for clock sources, where possible. They generally are the lower power choice. Shut down all peripherals not in use, like the Pulse Width Modulator (PWM), ADC, USART, etc. Use as many look-up tables as possible in your code, instead of using the CPU to compute the results. Check the power consumption of all external components. For instance, measure the voltage drop of all external resistors in the circuit. Lower the I/O pins that are used to power external peripherals such as serial EEPROMs or external analog devices. Another surprise can be your LEDs that are turned on. A single LED can wipe out your power-savings efforts. In general, look for current-consumption gremlins.

Conclusion

Device power savings in battery-powered applications are extremely important. You can achieve true value by using the microcontroller’s programmability. You can do this by changing the power-supply voltage at the output of a regulated charge pump. A second area would be to power down non-critical peripherals when not in use. Another option is to control the clocking strategy in order to optimize power versus functionality. Integrated circuit manufacturers are continuing to improve the dynamic performance of their peripheral devices while reducing the quiescent-current and supply-voltage requirements. Microcontroller manufacturers are adding modes, such as idle and sleep, that save average power over long periods of time. The combination of lower-power peripherals and microcontroller modes enhances your chances of having a low-power, battery-powered solution.

Got your checklist? Now take all of these variables and put on your low-power “state of mind” hat. You, as the perceptive programmer/hardware expert, need to evaluate each one of your applications and every situation inside those applications, to look for the power-consumption gremlins. Good luck!

This article is excerpted from Bonnie Baker’s book, “A BAKER’S DOZEN: Real World Solutions to Real World Analog Design Problems,” published Spring 2005 by Elsevier.

Muxing Your Micro With Style

By Bonnie C. Baker, Microchip Technology Inc.

Connecting external peripherals into your controller or processor is easy enough. You connect a processor I/O pin to each external peripheral. You then have independent control of each peripheral. Common peripherals are push-button switches or LEDs (Light Emitting Diodes). In reality, your controller or processor can connect to a larger variety of peripherals than listed here.

The trick is to use a microcontroller or microprocessor that has enough pins to connect to all the required peripherals. Unfortunately, each added peripheral translates to an additional device pin, which usually means into a higher application circuit cost. If you limit the number of pins on the processor or controller to limit cost, you have to use discretion as you add peripherals. This conflict between cost of controller/processor and number of required peripherals can present a difficult design problem.

The clever designer can work around these limitations. Adding an external multiplexer is a good first defensive move. If you use an external multiplexer, you can dramatically reduce the pin-count requirements for your controller or processor. You might exclaim that your cost model for your circuit does not allow the addition of another device. This is a legitimate concern. So maybe this first suggestion is not the right solution for you.

Rather than playing defense, consider changing to the offensive team. An external multiplexer is still thinking inside the “box.” The box that I am referring to is a device-oriented box. In other words, the premise is that adding a function implies adding a device. Consider an offensive move in which you solve this problem with an “out-of-the-box” solution. Figure 1 illustrates an alternative using this type of design strategy.

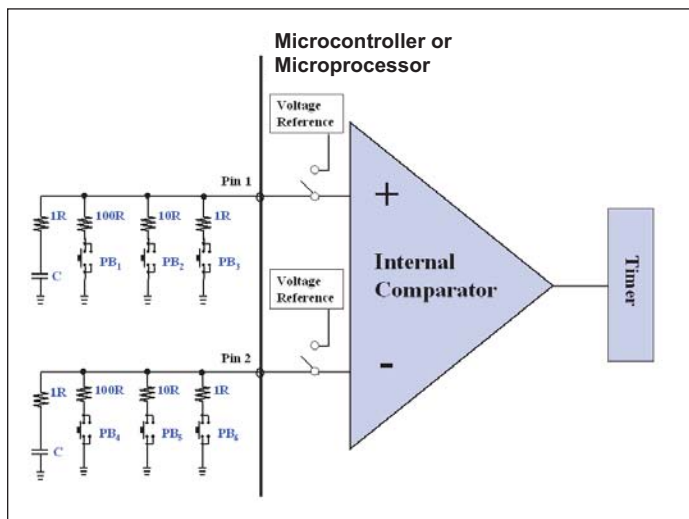


Figure 1: Two simple R/C networks connects six push buttons to a controller or processor. Pressing a push-button discharges the voltage from the capacitor through the resistor to ground. An alternative circuit connects each push-button to its own microcontroller or microprocessor pin.

In Figure 1 there are six push buttons attached to two ports of a microcontroller or microprocessor. This “outside the box” design approach reduces your required pin count from six down to two for the push-button functions. The circuit operates by first connecting pin 1 to the voltage reference. By biasing the voltage reference to its highest voltage, the capacitor, C, charges to its full positive value. After C is fully charged, the controller code then disconnects the voltage reference from the non-inverting input of the comparator. Unless the user presses one of the push buttons, the non-inverting input of the comparator will remain high during the duration of the test. If the user does press a button, the voltage at pin 1 discharges over time to zero. The time constant of this discharge is equal to $(1/((x+1)RC))$ for a single push-button, where x is the multiple of the resistors. At that time, the voltage reference at the inverting port voltage of the comparator (pin 2) is approximately $\frac{1}{4} V_{DD}$. During the discharge time, the timer on the output of the comparator counts out the time required to toggle the output of the comparator from high to low. By using a look-up table, you can identify the push button or combination of push buttons in use. Testing the status of PB1, PB2 or PB3 uses this algorithm. Alternatively, testing the status of PB4, PB5 and PB6 uses a similar algorithm.

Now imagine that you have a larger variety of peripherals in your application. With only two pins used, the other pins are available for additional peripherals. This multiplexing technique is portable to a variety of peripherals connected to the same pins. If your controller or processor has an internal comparator, voltage reference and timer, you can use resistors and capacitors as your external, low-cost “multiplexer”.

Miscellaneous Articles

Riding The Crest Factor Wave

By Bonnie C. Baker, Microchip Technology Inc.

Crest factor is a unit-less figure of merit that engineers use to describe a signal. You can use crest factors to describe the purity of voltage, current or power waveforms. This figure of merit can describe power signals, machine vibration, or half-wave lamp dimmers, to name a few. In these applications, it is not enough just to know the average power (or voltage or current) but you also need to know the equipment-design, peak-value guidelines in order to accommodate maximum allowable excursions. Within these disciplines, the crest factor is equal to the ratio of the peak value to the rms value. (Crest Factor = $V(\text{peak}) / V(\text{rms})$; $V(\text{rms})$ is the statistical standard deviation of the signal.)

A perfect sine wave with an amplitude peak of one volt has an rms value of 0.707V or $1/\sqrt{2}$ V. The crest factor for this signal is equal to 1.414 or $\sqrt{2}$. A sine wave that has a crest factor greater than 1.414 is less than perfect, showing additional noise, distortion or spikes. Unfortunately, non-periodic events will not appear in a Fast Fourier Transform (FFT). The crest factor will be higher than ideal if low-level noise or spikes exist.

There is another way to use crest factors in your system if you are not sure about the magnitude of the signal but would like to know the signal boundaries over time. You can take the above formula and rewrite it to solve for the signal peak value; $V(\text{peak}) = \text{Crest factor} * V(\text{rms})$ or $V(\text{peak-to-peak}) = 2 * \text{Crest factor} * V(\text{rms})$. This formula structure is extremely useful if you have a general idea about the type of signal that you are working with. For instance, crest factors are well defined for signals that form a gaussian or normal distribution across multiple samples. A gaussian distribution of noise is generated by analog systems that are void of internal switching circuitry. The common elements that generate gaussian noise are amplifier circuits, resistors, voltage references, D/A converters and ADCs. These devices generate non-periodic events riding on top of the signal of interest. Again, an FFT analysis will not produce an accurate indication of analog system noise.

When you compute the rms value of your system you must take multiple, periodic samples of a DC signal. Once you have these samples, you can calculate the rms value (see Figure 1). At this point, you need to determine how often you will accept an occurrence that exceeds the expected output value (see Table 1).

From the list of devices above, the sigma-delta ADC holds my interest the most. This is because the sigma-delta converter replaces an entire discrete analog system. The sigma-delta converter oversamples an analog signal and produces a stream of single bits. Then the converter employs digital filtering to calculate the one-bit conversion results to a higher resolution. In this process, gaussian noise is a by-product. The range of the converter's resolution change can be from one bit to 24 bits (or more). The accuracy (which is different than resolution, Ref 1) can be higher than 22 bits (rms) in this type of system. From these "noisy" results, you need to determine the acceptable system accuracy. With a crest factor, you are armed to effectively estimate the accuracy of your system (see Table 1).

Crest Factor (CF)	Crest Factor in Bits (BCF, bits)	Percentage of Occurrences Where Peaks are Exceeded
2.6	2.38	1%
3.3	2.72	0.1%
3.9	2.94	0.01%
4.4	3.13	0.001%
4.9	3.29	0.0001%

Table 1: Noise volts (peak-to-peak) in a signal is equal to Noise volts (rms) * 2 * CF. The Noise bits (peak-to-peak) is equal to Noise in bits (rms) – BCF. From the selected crest factor, you can predict the probability of an occurrence that exceeds defined peak-to-peak limits.

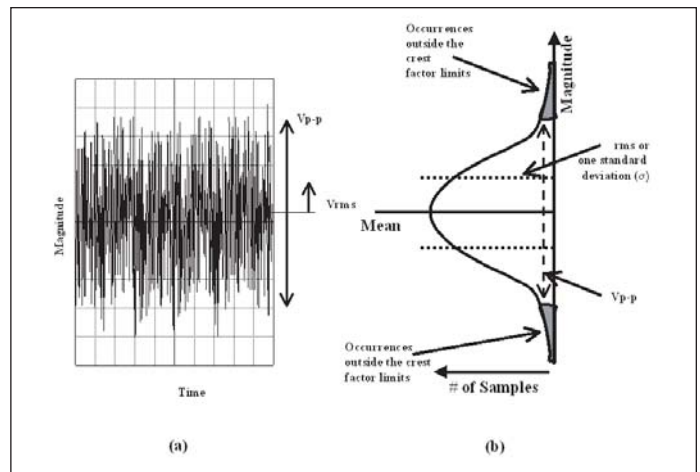


Figure 1: The standard deviation (rms) of several samples (>1024) of a DC signal is enough data to describe the system's future performance (per Table 1).

Designing With Low Dropout Regulators In Embedded Applications

By Qi Deng, Microchip Technology Inc.

Low dropout regulators (LDOs) are a mainstay of embedded microcontroller based applications. For designs where *simplicity*, *flexibility*, *footprint* and *cost* concerns are primary, LDOs are still the power sources of choice. These target applications include industrial and enterprise applications such as industrial controls, switch mode power supply (SMPS) post regulators, server boards and consumer products such as video/audio systems and Set Top Boxes (STBs).

A common characteristic of embedded applications is they are AC powered by either high voltage multiphase industrial or 110V/220V 50 Hz/60 Hz consumer AC sources. The AC voltage is then rectified to a high DC voltage such as 12V or 24V, which is further converted down to a low DC voltage. This DC/DC conversion is typically done by a step-down switching regulator to maximize power conversion efficiency. The low DC voltage, with 3.3V and 5.0V as the most common values, is the top level DC power source, or “system rail,” for the embedded circuitry. A majority of embedded microcontrollers these days are powered by a DC voltage ranging between 1.8V and 5.0V, which is typically converted from the “system rail” by a cost effective LDO device.

In addition to executing application level functions, the embedded microcontroller also acts as a “system supervisor,” performing tasks such as power condition monitoring and power management. Because of this, the LDO also needs to provide relevant power condition information to the embedded microcontroller to help it make system control decisions, and be controllable by the embedded microcontroller to react to system status changes.

Therefore, in addition to the *load current*, which is always the first parameter for designer to consider, several LDO specifications and features are considered critical in an embedded application, including *dropout voltage*, a power condition indicator output to the microcontroller, such as *power good* output, and a control feedback feature, such as *shutdown*.

It is always a good design practice to choose a LDO with the smallest dropout voltage for it offers the highest flexibility and fault tolerance. For example, in many embedded applications, a 3.3V system rail is converted down to 2.5V, leaving only 800 mV maximum dropout headroom over a wide temperature range (e.g., -40°C to 125°C). This means the typical dropout voltage under a normal condition (25°C) cannot be much more than 300 mV.

Power good indicator output is also critical because the embedded microcontroller needs to perform certain system housekeeping functions as a direct response to power level fluctuation. In an embedded application that converts a 3.3V system rail to 2.5V, if the system rail drops to, say, 2.8V, there is not enough dropout headroom to keep the regulation at 2.5V. When such a condition occurs, the power good output is asserted to notify the embedded microcontroller. Because the most likely response of the embedded microcontroller is to reset the system, a *programmable delay* on the power good output is desirable to optimize time setting of the system reset.

Another critical LDO feature is the need to be shut down by either the embedded microcontroller or some other system level intelligence. This is because that when the embedded microcontroller decides to switch off the system either voluntarily or as a response to fault conditions, the LDO needs to be shut down to preserve power or prevent damage.

Putting this all together, an ideal LDO for a typical embedded application should have the following critical features included in a single package:

- High load current, up to 1A
- Low dropout voltage, less than 300 mV typical
- Power good output with programmable delay
- Shutdown input

One example of an ideal LDO solution is the MCP1726 from Microchip Technology. The MCP1726 is the first LDO in the industry specifically designed for embedded applications with all the above features incorporated on one chip, including 1A load current, 150 mV typical dropout voltage, Power Good output with User Programmable Delay and Shutdown.

In addition, the MCP1726 offers other desirable features for numerous applications, such as small and thermal capable DFN (3 mm x 3 mm) package, stability with 1 μ F ceramic output capacitor, low 140 μ A typical supply current and low 0.001 μ A typical leakage current during shutdown. The standard output voltages for the MCP1726 include 0.8V, 1.2V, 1.8V, 2.5V, 3.3V, 5.0V, as well as adjustable, covering a complete input voltage range of embedded microcontroller. The MCP1726 is also available in more traditional SOIC package.

With the MCP1726, system designers no longer have to use up to three devices to perform all necessary power monitoring, conversion and management tasks. The simplicity, flexibility, footprint and cost advantages of the MCP1726 make it the top choice for embedded applications.

For more information on LDOs or the MCP1726, please visit Microchip's web site at: www.microchip.com

Analog Design Notes

Notes:

Notes:

Analog Design Notes

Analog and Interface Attributes

Robustness

- MOSFET Drivers lead the industry in latch-up immunity/stability

Low Power/Low Voltage

- Op Amp family with the lowest power for a given gain bandwidth
- 600 nA/1.4V/10 kHz bandwidth Op Amps
- 1.8V charge pumps and comparators
- Lowest power 12-bit ADC in a SOT-23 package

Integration

- One of the first to market with integrated LDO with Reset and Fan Controller with temperature sensor
- PGA integrates MUX, resistive ladder, gain switches, high-performance amplifier, SPI™ interface

Space Savings

- Resets and LDOs in SC70, A/D converters in a 5-lead SOT-23 package
- CAN and IrDA® Standard protocol stack embedded in an 18-pin package

Accuracy

- Offset trimmed after packaging using Non-Volatile Memory

Innovation

- Low pin-count embedded IrDA Standard stack, FanSense™ technology
- SelectMode™ operation

For more information, visit the Microchip web site at: www.microchip.com

Analog and Interface Products

Thermal Management

- Temperature Sensors
- Fan Speed Controllers/ Fan Fault Detectors

Power Management

- LDO & Switching Regulators
- Charge Pump DC/DC Converters
- Power MOSFET Drivers
- PWM Controllers
- System Supervisors
- Voltage Detectors
- Voltage References

Battery Management

- Li-Ion/Li-Polymer Battery Chargers
- Smart Battery Managers

Interface

- CAN Peripherals
- Infrared Peripherals
- LIN Transceiver
- Serial Peripherals

Mixed-Signal

- A/D Converter Families
- Digital Potentiometers
- System D/A Converters
- V/F and F/V Converters

Linear

- Op Amps
- Programmable Gain Amplifiers
- Comparators
- Linear Integrated Devices

Worldwide Sales and Service

At Microchip, we know that it takes more than product specifications to create loyal customers. In addition to a broad product portfolio, we understand the value of a complete design solution. That's why we maintain a worldwide network of sales and support. Our global network of experienced field application engineers and technical support personnel are ready to provide product and system assistance to help you further streamline your design, prototype and production activities.

Sales Office Listing

Technical Support: <http://support.microchip.com>

AMERICAS

Atlanta

Tel: 770-640-0034

Boston

Tel: 774-760-0087

Chicago

Tel: 630-285-0071

Dallas

Tel: 972-818-7423

Detroit

Tel: 248-538-2250

Kokomo

Tel: 765-864-8360

Los Angeles

Tel: 949-462-9523

San Jose

Tel: 650-215-1444

Toronto

Mississauga, Ontario

Tel: 905-673-0699

ASIA/PACIFIC

Australia - Sydney

Tel: 61-2-9868-6733

China - Beijing

Tel: 86-10-8528-2100

China - Chengdu

Tel: 86-28-8676-6200

China - Fuzhou

Tel: 86-591-8750-3506

China - Hong Kong SAR

Tel: 852-2401-1200

China - Qingdao

Tel: 86-532-8502-7355

China - Shanghai

Tel: 86-21-5407-5533

China - Shenyang

Tel: 86-24-2334-2829

China - Shenzhen

Tel: 86-755-8203-2660

China - Shunde

Tel: 86-757-2839-5507

China - Wuhan

Tel: 86-27-5980-5300

China - Xian

Tel: 86-29-8833-7250

ASIA/PACIFIC

India - Bangalore

Tel: 91-80-2229-0061

India - New Delhi

Tel: 91-11-5160-8631

India - Pune

Tel: 91-20-2566-1512

Japan - Yokohama

Tel: 81-45-471-6166

Korea - Gumi

Tel: 82-54-473-4301

Korea - Seoul

Tel: 82-2-554-7200

Malaysia - Penang

Tel: 604-646-8870

Philippines - Manila

Tel: 632-634-9065

Singapore

Tel: 65-6334-8870

Taiwan - Hsin Chu

Tel: 886-3-572-9526

Taiwan - Kaohsiung

Tel: 886-7-536-4818

Taiwan - Taipei

Tel: 886-2-2500-6610

Thailand - Bangkok

Tel: 66-2-694-1351

EUROPE

Austria - Weis

Tel: 43-7242-2244-399

Denmark - Copenhagen

Tel: 45-4450-2828

France - Paris

Tel: 33-1-69-53-63-20

Germany - Munich

Tel: 49-89-627-144-0

Italy - Milan

Tel: 39-0331-742611

Netherlands - Drunen

Tel: 31-416-690399

Spain - Madrid

Tel: 34-91-352-30-52

UK - Wokingham

Tel: 44-118-921-5869



MICROCHIP
www.microchip.com

Microchip Technology Inc. • 2355 W. Chandler Blvd. • Chandler, AZ 85224-6199

Microcontrollers • Digital Signal Controllers • Analog • Serial EEPROMs

The Microchip name and logo, the Microchip logo, Accuron, dsPIC, KeeLoq, microID, MPLAB, PIC, PICmicro, PICSTART, PRO MATE, PowerSmart, rPIC, and SmartShunt are registered trademarks of Microchip Technology Incorporated in the U.S.A. and other countries. AmLab, FilterLab, Migratable Memory, MXDEV, MXLAB, PICMASTER, SEEVAL, SmartSensor and The Embedded Control Solutions Company are registered trademarks of Microchip Technology Incorporated in the U.S.A. Analog-for-the-Digital Age, Application Maestro, dsPICDEM, dsPICDEM.net, dsPICworks, ECAN, ECONOMONITOR, FanSense, FlexROM, fuzzyLAB, In-Circuit Serial Programming, ICSR ICEPIC, Linear Active Thermistor, MPASM, MPLIB, MPLINK, MPSIM, PICkit, PICDEM, PICDEM.net, PICLAB, PICtail, PowerCal, PowerInfo, PowerMate, PowerTool, rLAB, rPICDEM, Select Mode, Smart Serial, SmartTel, Total Endurance and WiperLock are trademarks of Microchip Technology Incorporated in the U.S.A. and other countries. SQTP is a service mark of Microchip Technology Incorporated in the U.S.A. All other trademarks mentioned herein are property of their respective companies. © 2005, Microchip Technology Incorporated, All Rights Reserved. Printed in the U.S.A. 9/05

DS21975B

