Designing Electronics That Work (Sample Chapter) © 03/25/25 by Hunter Scott

4

SELECTING ACTIVE COMPONENTS

Generally, active components will define the major capabilities and specifications of your designs. You'll spend most of your component-selection time looking for your active components, so it's important to look for the right thing. In this chapter, we'll look at some of the most common categories for active components and discuss what to consider when selecting them.

Advertised vs. Actual Performance

As we explore active components, keep in mind that datasheets have been known to fib about the performance of some parts. After you gain some experience scanning distributor websites for particular parts, you'll get an idea of what's on the market. If you see a part that claims performance that seems too good to be true, it probably is. Datasheets love to use flashy claims on the first page to draw you in. For example, wideband amplifiers that claim 20 dB of gain will, upon closer inspection, actually only give you 20 dB of gain at exactly one input voltage, one input RF power, and one frequency.

Designing Electronics That Work (Sample Chapter) © 03/25/25 by Hunter Scott

For a more realistic view of a part's performance, test it on an evaluation board first, especially if you're planning to use the part at or near any of its limits. You can sometimes get evaluation boards for free through sales engineers, but you may need to buy them. If they're too expensive, you can try fabricating one yourself, since evaluation board Gerber and manufacturing files are almost always released publicly or available upon request. The idea is that you can verify that the datasheet is telling the truth and that the part will perform the way you expect in your application before you commit to ordering it in large quantities.

Oscillators and Crystals

When oscillators are used as an external reference for an IC, the datasheet and application notes often give oscillator parameters needed to guarantee nominal IC operation. They'll sometimes also recommend specific part numbers. Follow these recommendations and use the suggested part if possible. That said, sometimes the suggested part numbers will be end-of-life or unavailable, so don't just blindly use the recommended parts without checking that they're still active and in stock with an authorized distributor.

Crystal oscillators typically need two external capacitors, called *shunt capacitors*, to function correctly. You can get the value of these capacitors from the datasheet or you can calculate it. A good reference for these calculations is the *ST Oscillator Design Guide* (available as a free PDF from ST Microelectronics). That design guide contains a wealth of information about oscillators and crystals, so rather than re-create it here, I recommend giving it a look. In most circumstances, the best thing to do is use the values specified by the oscillator's datasheet. Always use NP0/C0G-type capacitors. You can optionally decrease the capacitance by 1 or 2 pF depending on the physical size of the capacitors you're using to compensate for the extra capacitance you'll get between the capacitor terminals and PCB's ground plane. As long as you're within a couple of pF of the correct value, the oscillator should behave just fine.

If you plan on using a crystal oscillator outside of its temperature range, you need to characterize its performance yourself. The datasheet will only guarantee performance within a certain temperature range, and outside of that range, the crystal may experience an *activity dip*. This is the point where the crystal suddenly increases in resistance, which can cause a frequency error of up to 20 ppm that can make your system lose phase lock. These dips happen at very narrow temperature ranges, and they're almost never specified in the datasheet. If your design needs a highly accurate clock reference over a large temperature range, you should characterize the crystal you're planning to use and design around those points if necessary.

A crystal's resonant frequency also changes slightly over time as the crystal ages. Luckily, this change is very small, usually about 5 ppm during the first year and 10 to 20 ppm over the next 10 years. Still, for very sensitive designs, it's important to be aware of this fact. The aging effect isn't impacted by temperature.

Designing Electronics That Work (Sample Chapter) © 03/25/25 by Hunter Scott

Silicon oscillators offer several advantages over crystal oscillators. Instead of a piece of quartz vibrating, it's a MEMS structure on silicon. This allows silicon oscillators to be physically smaller than crystal oscillators and allows them to be programmable to a range of frequencies. Silicon oscillators can operate from 1 kHz up to 170 MHz and include EMI-reduction features.

Before 2018, silicon oscillators could fail if exposed to even small concentrations of helium gas. There was an incident where a bunch of technicians installing an MRI machine had their iPhone 8s all die because of a small helium leak. Eventually, the gas dissipated and the phones started working again. It turned out that helium molecules were small enough to migrate inside the silicon package of the oscillator and cause a significant frequency deviation. Since then, manufacturers have developed special packages and coatings that reduce or eliminate this problem. If you decide to use a silicon oscillator, you may want to find one that's rated to operate in atmospheres containing concentrations of small gas molecules, especially if you expect your product to have to operate under those conditions.

ADCs and DACs

There's a huge amount of information to consider when choosing an ADC or DAC, and there simply isn't room in this book for the discussion they deserve. Instead, check out TI's "Choose the Right A/D Converter for Your Application" and *The Data Conversion Handbook* by Analog Devices (2005). They're both linked on the website for this book.

Power Supplies

When choosing power supplies, take the time to calculate your power budget by determining the minimum and maximum current and power that each component in your design will require. Then make sure your power supplies can provide at least 10 percent more than what you need. Also make sure none of your components are dissipating more power than they're rated for, and that you've taken measures to handle any heat produced by those parts.

Voltage Conversion

A switched-mode power supply or DC/DC converter is the best way to go from a low voltage to a high voltage and vice versa. This is useful when you need to move large amounts of power. It's more efficient to send the power using a high voltage; the current will be lower, which will reduce your loss due to resistance in your traces, cables, or transmission lines. However, since most designs tend to run at a low voltage, you'll need to reduce the voltage back down again, which is where an SMPS comes in.

There are myriad SMPS and DC/DC converter ICs that you can use. TI has a nice, free online tool called WEBENCH (*https://webench.ti.com/*) that

Designing Electronics That Work (Sample Chapter) © 03/25/25 by Hunter Scott

will help you run some simulations and find a part that meets your requirements. SMPSs can be fickle if you don't use the right inductor, use a poor layout, don't have enough output capacitance, or run the part outside of its operating range. Make sure to use as many of the suggested parts in the application note as possible and follow the layout guidelines. There's more information about that in Chapter 7.

If the FET in an SMPS is switching below 20 kHz, you'll hear a soft humming or squealing noise coming from the inductor and capacitors. This is because some capacitors are slightly piezoelectric and some wire-wound inductors will physically move slightly as their magnetic field changes, creating sound. You can solve this problem by running the SMPS above 20 kHz. This way, any sound produced will be above the human hearing range.

Sound shouldn't be the only reason you pick a particular switching frequency, though. More importantly, the switching frequency decides what value inductor you need, and the higher the inductance you need, the larger the physical inductor must be. Because there are so many different SMPS ICs, topologies, and applications, discussing exactly how to choose a switching frequency for your design is out of the scope of this book. The datasheet for your SMPS IC should have a section on switching frequency and usually some suggested inductors. The TI WEBENCH tool will also let you sort SMPS ICs by the physical layout area required for operation, which includes the inductor size. It's common for chips to abstract away the switching frequency design decision for you and simply call out a single inductor value to use to get the performance described in the datasheet. In any case, you need to be aware of the size of the inductor required for a particular IC, especially if you have a space-constrained design. Semiconductor companies can get tricky on you because they'll have big, glossy advertisements about how small their SMPS IC package is, but conveniently leave out that the size of the required inductor is just as big or bigger.

SMPSs can cause your power lines to have extra voltage ripple caused by the inductor switching on and off. This ripple is characterized in the datasheet of the IC, but it's critical that you use the right capacitors to achieve datasheet performance. The number of capacitors, their type, where they're placed, and their value all determine the cleanliness of the power on the output of the SMPS. If ripple is important to you, check the part's datasheet before you choose to use it, and make sure it's acceptable even after derating it slightly.

Make sure you use DC/DC converters exactly as described in their datasheets and application notes. If you're using a DC/DC converter in a nonstandard way, talk to the application engineers or simulate it using a model from the manufacturer and real (non-ideal) components. LTspice is a good simulation tool for those kinds of situations, and it's free. Keep in mind that DC/DC converters are less stable at small (less than 20 percent) loads. Also, don't use the synchronization feature on DC/DC converters unless you really need it. It's easy to cause increased noise or oscillation if the feature isn't used exactly right.

Designing Electronics That Work (Sample Chapter) © 03/25/25 by Hunter Scott

Following the application notes for DC/DC converters will get you to a working circuit, but applications notes don't always show the most optimal way to design with the part. They're usually pretty good, but, as an example, some DC/DC converter application notes will recommend placing a snubber circuit near the switching transistor to help reduce EMI. This will reduce EMI, but it will also reduce the DC/DC converter's efficiency. As a rule of thumb, follow the application note, but keep your wits about you and know that it's okay to change things if you know what you're doing.

Low Dropout Regulators

If you use low dropout regulators (LDOs), consider using parts rated for low noise. They usually don't affect cost and can have significantly lower ripple on the output. They can also save you a few components by reducing the number of output capacitors you may need.

Remember to compensate for the dropout voltage of your LDOs in your power budget. If you want 5 V on the output, you need to provide *more* than 5 V on the input. The lower the dropout of the part, the better, since less power will be lost as heat. LDOs operate most efficiently when their input voltage is *just* high enough to account for the drop, since any higher voltage will just be burned out as heat.

LDOs have a *power supply rejection ratio (PSRR)* that's given in the datasheet. This is the attenuation in dB of the noise on the input of the LDO compared to the output. It's not uncommon for this value to be 40 to 50 dB, or even higher, depending on the dropout voltage. You can take advantage of this by using LDOs as a power supply filter. This can be helpful for removing ripple from an SMPS or any noise picked up by a long power cable. A common technique is to use an SMPS to do large voltage conversions more efficiently, and then put an LDO with a low dropout right before the ICs you're trying to power to help isolate them from the noise of the rest of the power nets.

Thermal Shutdown

One nice feature that many power supply ICs contain is thermal shutdown. An internal temperature sensor will automatically disable the regulated output once the chip gets too hot. Thermal shutdown is often triggered at very high temperatures, so reliability of the part will be poor if it's operated in a way that triggers thermal shutdown with any regularity. This feature is meant more as a safety measure.

Power supply ICs will sometimes also have a "power good" pin, which you can use to drive an LED or check with a microcontroller to detect fault conditions and take the correct action. Look for both of these features when shopping for power supply chips.

Designing Electronics That Work (Sample Chapter) © 03/25/25 by Hunter Scott

Microcontrollers

Before you go all in on a microcontroller, buy a development board for it. This will give you a good idea of how painful it is to actually get code running on the chip, and you'll get a feel for what development will be like. Development boards sold directly by the manufacturer are sometimes very expensive, but you can almost always find a cheaper third-party development board. Check sites like Tindie, eBay, or AliExpress. I also recommend talking to someone who has used the chip you're considering. They can be a great help in getting past the marketing veneer and finding out what it's really like to develop on the chip.

Before you decide on a microcontroller or processor, you should read its errata. This is a separate document you need to download from the manufacturer's website that lists all of the known problems or mistakes in the silicon, what their symptoms are, and what (if any) workarounds exist. Reading this document can save you days or weeks of troubleshooting and, in extreme cases, can cause you to pick a different part.

Meeting Your Requirements

Look for microcontrollers that already contain as many of your requirements as possible. For instance, if your requirements call for Ethernet, try to find a microcontroller with built-in Ethernet instead of having to use an external Ethernet IC. Avoid as much bit-banging as you can. A hardware implementation of a bus or protocol on your microcontroller will save you lots of development time, will generally have better performance, and is less likely to have bugs. On the other hand, bit-banging can be a great way to work around a buggy hardware peripheral.

Besides Ethernet, other common buses that some microcontrollers include are USB, I2C, SPI, USART, and UART. Different chips have different numbers of each of those buses, so if you have a UART-heavy application, for example, try to find a chip that includes multiple UARTs. You can also find microcontrollers that support specialized protocols like MIDI or CAN bus, and there are even microcontrollers with built-in radio modules. This can save you a lot on your BOM, since you can eliminate an entire external IC and just use your microcontroller. Sometimes microcontrollers won't let you use all peripherals all at the same time because of limitations in pin configurations, so you should download the manufacturer's configuration tool to verify that you can actually use the combination of peripherals you want.

When it comes to memory requirements, estimate the amount of both RAM and nonvolatile storage that you'll need. If you need more nonvolatile storage than your preferred microcontroller can provide, you can use an external storage component. SPI EEPROMs are common and can store huge amounts of information (this can be useful for storing lookup tables to improve application performance). Talk to your software team about how much RAM and storage they'll need to meet the design requirements. You'll also need them to weigh in on performance requirements: the number of cores and threads, clock speeds, and whether any cryptographic or graphics acceleration cores are needed. If your software team is already familiar with a particular processor family or architecture, or has a lot of code and boilerplate written for it, you should probably give more weight to that processor family or architecture when selecting one.

Considering Software Development Environments

You should also investigate the software development environment required for a microcontroller before deciding on a part. If you decide to use the proprietary tools that the manufacturer provides, test them out first and take a look through the provided example code and libraries. You may be surprised at the poor quality of the code in the provided libraries, and it's likely that there will be bugs.

If you decide to use the manufacturer-provided IDE, make sure to budget for it. Sometimes they can be extremely expensive. Other times they cost money but have a free, feature-reduced version. For example, IAR Workbench has a free version that only lets you compile up to 32 KB of object code. Not realizing this beforehand could cause you to exceed your budget or waste time optimizing code size if you can't afford the IDE.

You also have the option of using a free, open source tool chain. There are lots of tools and support for almost every microcontroller out there, but because of that, you'll probably need to take some time to get your development environment set up for your particular chip. One IDE that does a good job of integrating a lot of tools together for you is PlatformIO (*http://platformio.org*). It doesn't have support for every single IC, but it does support a lot, especially for IoT applications.

Derating the Operating Parameters

It's always a good idea to derate microcontroller operating parameters so you aren't running them at their maximum ratings. Table 4-1 shows NASA's recommendations for derating integrated circuits.

Table 4-1: Integrated Circuit Derating

Stress parameter	Derating (digital)	Derating (linear)
Max supply/input voltage	90%	80%
Power dissipation	80%	75%
Max junction temperature	80%	75%
Max output current	80%	80%
Max clock frequency	80%	80%

Keep in mind that these are guidelines, and since these guidelines are from NASA and intended for spacecraft design, they're conservative.

RF Components

There are a lot of things to consider when choosing RF components, and a full treatment of how to design an RF signal path is out of the scope of this book. In this section, we'll look at a few of the more universal RF part selection considerations that apply to complex designs as well as simpler and more commonly encountered RF circuits.

Frequency-Dependent Parameters

If you're picking out an RF part, pay attention to the insertion loss at the frequency you'll be using it. Insertion loss changes based on frequency, but the datasheet will usually advertise the minimum insertion loss across its entire frequency operating range. If you're not careful, your link budget will be wrong because you didn't account for the extra loss at your actual frequency of operation. As always, you need to check the plots in the datasheet and not go by the bullet points on the first page.

The higher the power that you're working with, the more important loss is in your parts. Consider a cable that you measure to have 1 dB of loss. At a low power, the difference of a single dB is a small amount of absolute power. But at a high power, a single dB can be watts of difference. This is why it's usually a good idea to put the high power/gain stage right before your antenna. Even low-loss cables can cause you to drop significant power if they're too long or if you put them in the wrong place.

Gain is another frequency-dependent parameter that can be advertised in a misleading way. This is especially true for wideband amplifiers. No matter what the datasheet says, an amplifier can't be both perfectly wideband and perfectly high-gain. There will be a curve showing gain over frequency, and *that's* the chart you should look at, not the gain in huge print on the first datasheet page or in the ad. The gain is usually lowest near the edges of the advertised frequency range. For example, if a part claims that it has 30 dB of gain and works from 1 to 2 GHz, you probably won't see 30 dB of gain at 2 GHz. Very high-gain amplifiers will typically have a relatively small frequency range.

Amplifiers

If you want to amplify a signal up to any reasonably high power, you'll almost certainly need to use multiple amplifiers, or at least a single multi-stage amplifier. As such, don't be surprised if you can't find a part that will do, say, +30 dB of gain at +36 dBm out. Instead, maybe look for one amplifier that does +25 dB of gain at +26 dBm out, and a second amplifier that does +10 dB of gain at +36 dBm out. There's a tradeoff between the maximum power out of an amplifier and the maximum gain it exhibits.

Don't count on being able to achieve the maximum power out advertised for an amplifier, however, especially if it's a lot of power. Give yourself at least 1 or 2 dB of headroom, even if you're going to be running the amplifier into hard saturation. As always, the best thing to do is to get an evaluation board of the amplifier you're interested in and drive it the way you'll be driving it in your product to see what performance you get. These development boards are often several hundred dollars, but you can usually get them for free from a sales engineer. You can also check eBay or AliExpress for unofficial development boards for some parts, although these usually don't come with documentation. Manufacturers will also usually give you the Gerbers for the development board for free so you can fabricate and assemble it yourself. The only downside with that approach is that you don't get a piece of paper with measured, verified performance data of your development board from the manufacturer.

In an ideal amplifier, gain is linear. If you make the input signal a little bit bigger, the output will be proportionally bigger. But real amplifiers only behave that way up to a point, after which the amplifier starts to go into compression. The input power where you no longer get a linear gain, but instead get a gain 1 dB less than you would expect an ideal part to have, is called the *1 dB compression point*, or P_{1dB}. This is an important number to look at, especially if you're going to be running an amplifier in compression or at its maximum output power. Figure 4-1 illustrates the difference between an amplifier's ideal response and its actual response and shows the $\mathrm{P_{1dB}}$ point.

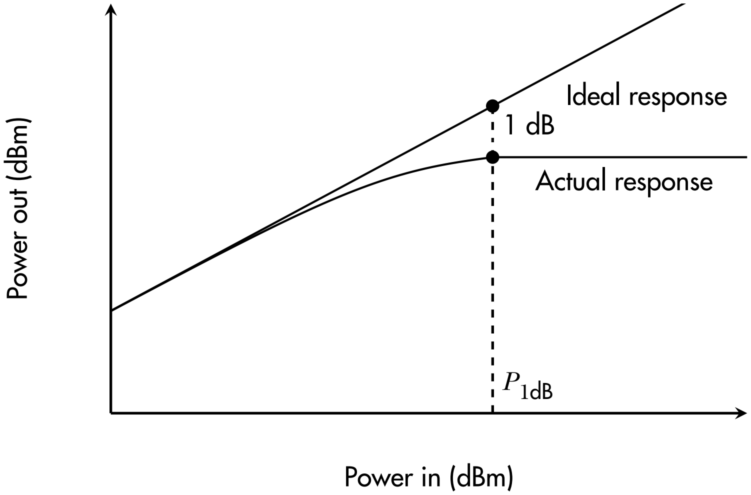

Figure 4-1: P_{1dB} is the point where an amplifier's gain is 1 dB less than it would ideally be.

Amplifiers also have a stability factor, or *k-factor*, that tells you if the amplifier will begin oscillating under poor impedance-matching conditions. You can find amplifiers that are both conditionally and unconditionally stable. Prefer amplifiers that are unconditionally stable. Unconditional stability means that no matter what impedance you put on the input or output of the amplifier, it won't oscillate. That doesn't mean it will *never* oscillate (there are other conditions you can subject it to that can cause oscillation), but using an unconditionally stable amplifier can save you a lot of troubleshooting.

Why would anyone ever want to use an amplifier that *isn't* unconditionally stable? As in everything, there are trade-offs in making an amplifier unconditionally stable. Specifically, the noise figure or P1dB may be worse in an unconditionally stable design versus a conditionally stable design. If these are trade-offs you'd like to play with, then you can choose a conditionally stable amplifier.

NOTE *If you want to read more about amplifier stability, there's a classic paper called "Avoiding RF Oscillation" by Les Besser (*Applied Microwave and Wireless, *spring 1995) with lots of good information. It's linked on this book's website,* https://designingelectronics.com.

Transistors

Discrete transistors find uses in devices like motor controllers, load switches, and power supplies. To use them successfully, there are a few factors you need to consider.

MOSFETs

MOSFETs should always be derated. Follow the NASA recommendations shown in Table 4-2.

Table 4-2: MOSFET Derating Guidelines

Stress parameter	Derating
Power	60%
Current	75%
Voltage	75%
Junction temperature	80%
Gate-to-source voltage	60%
Source-to-drain voltage	75%

MOSFETs don't instantly turn on and off. They have a gate capacitance that can be quite large in some cases. This capacitance adds a time delay before the MOSFET turns on. When picking a MOSFET, check that the gate capacitance isn't so large that you won't be able to meet timing requirements. The equation for the time it takes for a MOSFET gate to charge to the gate voltage and turn on is:

$$t = R_G \times C_{iss} \times \ln \frac{V_{GS} - V_{TH}}{V_{GS} - V_{gp}}$$

In this equation, R_G is the gate resistance, C_{iss} is the gate capacitance of the MOSFET as seen by the gate driving circuit at V_{DS}, V_{GS} is the final gate voltage, V_{TH} is the gate turn-on threshold voltage, and V_{gp} is the gate plateau voltage. You can obtain all of these values except V_{gp} from tables in the datasheet. To get V_{gp}, find the gate charge curve and determine the voltage where the curve plateaus. The gate charge curve will have V_{GS} on the y-axis and Q_g on the x-axis. Figure 4-2 shows a typical curve.

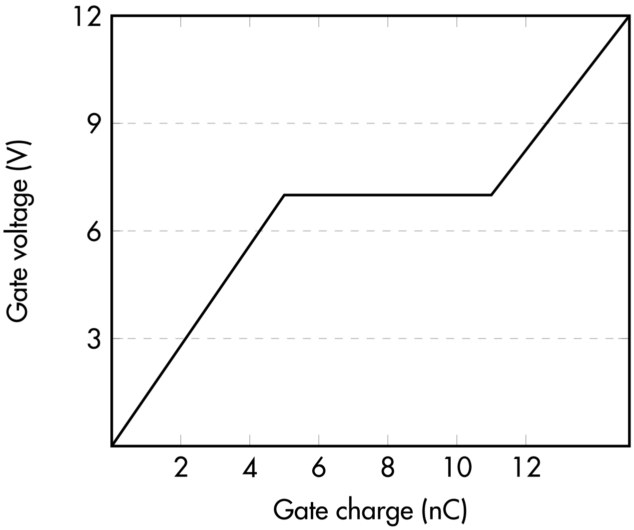

Figure 4-2: Example gate charge curve

In the example shown in the figure, the plateau occurs in the middle of the curve, at about 7 V. That's the value of V_{gp}.

NOTE *The full derivation of the gate charge time equation can be found in Vishay application note AN608A, available at* https://www.vishay.com/docs/73217/an608a.pdf.

Drivers

Determine if you need a driver IC or circuit to control any switching semiconductors. For example, a microcontroller or processor GPIO pin won't be able to provide enough current or voltage to fully switch a large power transistor. GaN FETs require a specific power-up sequence or else they'll break, so you need a special driving circuit for those devices. When choosing the driving method, pay attention to the switching times so that you can meet timing requirements. Having a fast MOSFET doesn't do you any good if you use a slow driver.

Drain-to-Source Resistance

All diodes and FETs have power dissipation because there's a nonzero drain-to-source resistance, or R_{ds}. If you're designing low-power or battery-operated electronics, a high R_{ds} will dissipate power unnecessarily. This can cause your measurement of consumed power to be higher than the value you calculated. In high-power devices that are switching large currents, a high R_{ds}

can cause your MOSFETs to overheat and fail prematurely or require significant thermal management.

Avoid these problems in the first place by checking the drain-to-source resistance during component selection. You can calculate how much power (in watts) your MOSFET will dissipate as a result of R_{ds} by simply squaring I_{ds} and multiplying it by R_{ds}. It's typically pretty easy to find a MOSFET that has an R_{ds} of under 1 Ω, so that's a good place to start.

It's important to note that R_{ds} increases as temperature increases, so if you're running your circuit in a hot environment or if you have insufficient heat sinking, your MOSFET will get hot, which will cause R_{ds} to increase. This will dissipate more energy as heat, which will make R_{ds} increase even more. You'll get stuck in thermal runaway until your MOSFET explodes. To avoid this, make sure that you calculate the expected thermal dissipation due to R_{ds} across your entire temperature specification.

Diodes

Choose LEDs to be sufficiently bright. LEDs should have a maximum forward current of less than the maximum current source and sink rating on the LED driver or GPIO pin you're using to drive them. If they don't, they should be buffered. You can't run an ultra-bright 10 W LED off of a microcontroller pin.

When using a diode in series, make sure you compensate for the voltage drop across the diode. If you need a really small forward voltage drop across a diode (usually the case when you use a diode in series), use a Schottky diode. The trade-off is that they have a higher reverse bias leakage current, and this gets worse as the Schottky diode's temperature increases.

Make sure the diode you pick can handle the current you need to put through it. You also need to specify the reverse voltage (the point at which the diode will begin to break down and start conducting in the "wrong" direction). Once you start approaching the rated reverse voltage, the diode will start to get leaky, and small amounts of current will start to flow. Use the I-V curve in the datasheet of the part you pick to figure out exactly how much current that is. If you don't give yourself enough headroom between the rated values and what you're actually subjecting the part to, you can end up with an inefficient design, since you'll be losing energy through diode leakage.

It's possible to use a Zener diode as a voltage regulator, but it's not recommended. A Zener diode regulator will have poor efficiency, poor power handling, and won't be very good at maintaining a regulated voltage. That said, they're very cheap, and if you have a low-power application that can handle those drawbacks, it may be an okay choice. A related application of Zeners that can work well is as a voltage reference. You can also use a Zener as part of a clamping circuit to ensure the voltage at a particular point never exceeds the reverse voltage of the Zener diode.

Be aware that low-voltage Zeners (below about 6 V) have a soft knee—that is, they start conducting before the breakdown point, making them look

Designing Electronics That Work (Sample Chapter) © 03/25/25 by Hunter Scott

"leaky." Zener diodes are also temperature sensitive. In fact, they have a negative temperature coefficient when operated below about 5 V and a positive temperature coefficient when operated above 5 V. The datasheet will usually contain a graph that shows you the temperature coefficient versus voltage.

Batteries

There are dozens of battery chemistries to choose from, and each one has its own advantages and limitations. Which chemistry you choose will depend on what fulfills your requirements, but 90 percent of designs can get by with one of the chemistries discussed in this section. If you're in the 10 percent that's unable to use any of these chemistries, you either need a more exotic chemistry (which is out of the scope of this book) or, more likely, need to try to change your requirements or specifications.

Remember that you need to properly dispose of all battery types. If you just throw batteries away, they can poison the environment at best, and explode or start a fire at worst.

If you need to measure detailed information about your battery performance, you can use a chip called a *gas gauge*. Gas gauges are all slightly different, but generally they have a Coulomb counter to measure charge accumulation into a battery, as well as current and voltage sensing. They can also sometimes provide digital outputs that report whether the battery voltage is sufficiently high, the charging state, and other useful pieces of information for managing battery charging or alerting your user.

Note that in battery parlance, a *primary battery* isn't rechargeable, while a *secondary battery* is rechargeable. A battery's capacity is measured in amp-hours (Ah) or milliamp hours (mAh). A battery with a rating of 1 Ah can discharge 1 amp for 1 hour, 0.5 amps for 2 hours, 2 amps for 0.5 hours, and so on.

Batteries also have a *C rating*, a dimensionless number used to describe the battery's charge and discharge rate. When a battery's C rating is listed, it's usually referring to the maximum discharge rate. For example, a 1 Ah battery with a 1 C rating means that it can discharge no more than 1 amp in 1 hour before damage starts to occur. The same 1 Ah battery with a 10 C rating means that it can discharge up to 10 amps for 0.1 hours (which is 6 minutes). However, just because you *can* discharge a battery at a huge C rate doesn't mean you should. Pulling a lot of current from a battery is going to make your battery get hot and will probably shorten its lifespan if you do it over and over. It's better to make sure you have a little room between the C rate you're using and the maximum C rating of your battery. Also, never charge batteries at a C rating above what the datasheet recommends. For most batteries, this is 1 C. Doing so could damage the battery.

Designing Electronics That Work (Sample Chapter) © 03/25/25 by Hunter Scott

Lithium

When people talk about lithium batteries, they're usually referring to secondary (rechargeable) lithium-ion batteries. Within the group of lithium-ion batteries, the two most common chemistries you'll encounter in consumer electronics are lithium polymer (LiPo) and lithium iron phosphate (LiFePO4).

In lithium polymer batteries, the electrolyte is a polymer instead of a liquid. These batteries have a high energy density and a nominal cell voltage of 3.7 V. The two most ubiquitous form factors of LiPo batteries are flat cells and 18650s. Flat cells are exactly what they sound like: thin, rectangular batteries. They come in all sizes, and cells of this kind are typically stacked on top of one another for multi-cell packs. Meanwhile, 18650 batteries have a standard-size cylindrical form factor. The name refers to the physical dimensions of the battery: 18 mm diameter by 65 mm long (just ignore the ending zero). There are other sizes of cylindrical cells available, but 18650 is by far the most common and popular. In addition to these two common form factors, LiPo batteries are also available as coin cells. There are many different sizes of coin cells, and you can buy both primary and secondary lithium coin cells.

LiPo batteries with multiple cells must be balanced. That means the voltage on each cell needs to be the same so they're all discharged at the same rate. Many multi-cell LiPo batteries come with balancing circuitry already included in the pack. Depending on how the battery is constructed, however, you'll probably still need to use a balancing external charger.

The failure mode you'll see if the batteries aren't balanced is that the pack will begin to get puffy as gas escapes and builds up inside. If a LiPo is puffy, stop using it. Puffy LiPos can heat up, self-ignite, and then self-oxidize. Because the battery creates its own oxidizer when burning, it's impossible to put out the fire with water. You just have to sit back and wait for the fire to burn itself out. Some battery packs have temperature probes built into them so the charger can detect any failure modes before they happen and stop charging. It's a good idea to use this feature if you can.

There are many cheap ICs designed explicitly for charging LiPo batteries, from one to many cells. This way, the battery can be charged while it's installed inside the device, rather than the user having to remove the battery and charge it externally. These charging ICs do more than just apply a voltage and balance the cells. They also change the voltage and current curves going into the battery to optimize its lifetime and charge it as quickly and safely as possible.

Lithium iron phosphate is a chemistry that has the advantages of LiPo while being much safer. When LiPo cells are shorted, they heat up and can cause a fire. LiFePO4 batteries won't catch on fire when punctured or crushed. The trade-off is that they have slightly less energy density than LiPo batteries and can cost slightly more.

Another new lithium chemistry that's now commercially available is called lithium thionyl chloride. These batteries are primary (not rechargeable), but they can have an extremely high energy density. A battery of this

Designing Electronics That Work (Sample Chapter) © 03/25/25 by Hunter Scott

chemistry the size of a D cell alkaline battery can hold up to 20,000 mAh at 3.6 V.

Regardless of the type, lithium batteries should be stored at around 40 percent charge to maximize their lifetime. Most consumer products that use lithium batteries will be packaged and shipped at about 40 percent to 70 percent charge. However, it's important not to allow the cells to discharge below 2.0 V per cell or the battery may be permanently damaged.

Lead Acid

Lead acid batteries are almost never used in consumer electronics, but they are sometimes used in robotics, automotive engineering, or other contexts where you need a very large energy density and a large current surge, but don't care much about battery size or weight. Lead acid batteries are very heavy (since they contain mostly lead) and have a cell voltage of 2 V. You usually see 6 V, 12 V, and 24 V lead acid battery packs.

Lead acid batteries are relatively safe. However, they do contain acid, and they also release hydrogen gas when they're charging, which can explode under the right circumstances. They're easy to recycle, provided you don't allow the lead to leach into the environment. To prolong their lifetime, lead acid batteries should be stored at 100 percent charge.

Alkaline

Alkaline batteries are primary batteries, meaning they're not rechargeable. This is the chemistry that AA, AAA, C, and D batteries use. Alkaline batteries are the cheapest way to implement battery power because they don't require recharging circuitry. However, they do require extra mechanical engineering, since your user needs a way to replace them when they die, whereas a rechargeable battery can be charged just by exposing a small connector somewhere on the device. Keep in mind that removable battery covers can cause problems with dust or water ingress into the device.

Correctly stored, a primary alkaline battery can have a shelf life of 10 years. Because alkaline batteries are primary cells, you don't need to discharge them at all before you store them. There will be some self-discharge over time, but it will be fairly minimal. Like lithium batteries, alkaline batteries are also available as coin cells (although not all coin cells are alkaline; some use different chemistries).

Conclusion

Active components are what give your product the features it needs. They're also probably the most expensive parts in your design. They consume the most power and get the hottest, too. Choosing and using your active components is a key part of the design process. In this chapter, we looked at all kinds of active components and what to consider when choosing them.