

- Strategy 1.* Focus on a specific component.
- Strategy 2.* Use well-studied experimental manipulations.
- Strategy 3.* Focus on large components.
- Strategy 4.* Isolate components with difference waves.
- Strategy 5.* Focus on components that are easily isolated.
- Strategy 6.* Use component-independent experimental designs.
- Strategy 7.* Hijack useful components from other domains.

### 3 Basic Principles of ERP Recording

This chapter describes how to record clean, artifact-free data. As the ERP technique has matured, there has been a decrease in the amount of discussion in the literature of basic issues such as recording clean data. This is only natural, because a number of laboratories have developed excellent techniques over the years, and these techniques have become a part of the laboratory culture and are passed along as new researchers are trained. However, as time passes, the reasons behind the techniques are often lost, and many new laboratories are using the ERP technique, making it important to revisit the basic technical issues from time to time.

#### The Importance of Clean Data

Before I begin discussing these issues, I want to discuss why it is important for you to spend considerable time and effort making sure that you are recording the cleanest possible data. The bottom line is that you want to obtain experimental effects that are replicable and statistically significant, and you are unlikely to obtain statistically significant results unless you have low levels of noise in your ERP waveforms. As discussed in chapter 1, the background EEG obscures the ERPs on individual trials, but the ERPs can be isolated from the EEG noise by signal averaging. As you average together more and more trials, the amount of residual EEG noise in the averages will become progressively smaller, so it is crucial to include a sufficient number of trials in your ERP averages. However, increasing the number of trials only works well up to a point, because the effect of averaging on noise is not a direct, linear function of the number of trials; instead, the noise decreases as a

function of the square root of the number of trials in the average. As a result, you can't cut the noise in half by doubling the number of trials. In fact, doubling the number of trials decreases the noise only about 30 percent, and you have to include four times as many trials to reduce the noise by 50 percent. Chapter 4 will cover this in more detail.

It should be obvious that you can quadruple the number of trials only so many times before your experiments become absurdly long, so increasing the number of trials is only one part of the solution. The other part is to reduce the noise before it is picked up by the electrodes. Much of the noise in an ERP recording arises not from the EEG, but from non-EEG biological signals such as skin potentials and from electrical noise sources in the environment accidentally picked up during the EEG recording, and it is possible to reduce these sources of noise directly. In fact, if you spend a few days tracking down and eliminating these sources of noise, the resulting improvement in your averaged ERPs could be equivalent to the effects of doubling the number of trials for each subject or the number of subjects in each experiment. This initial effort will be well rewarded in every experiment you conduct.

In addition to tracking down noise sources and eliminating them directly, it is possible to reduce noise by using data processing techniques such as filtering. As chapter 5 will discuss, these techniques are essential in ERP recordings. However, it is important not to depend too much on post-processing techniques to "clean up" a set of ERP data, because these techniques are effective only under limited conditions and because they almost always distort the data in significant ways. This leads us to an important principle that I call *Hansen's Axiom*:

*Hansen's Axiom:* There is no substitute for good data.

The name of this principle derives from Jon Hansen, who was a scientist and technical guru in Steve Hillyard's lab at UCSD. As he put it in the documentation for a set of artifact rejection procedures:

*There is no substitute for good data. It is folly to believe that artifact rejection is going to transform bad data into good data; it can reject occasional artifactual trials allowing good data to be better. There is no way that artifact rejection can compensate for a subject who consistently blinks in response to particular events of interest or who emits continuous high-amplitude alpha activity. In other words, data that are consistently noisy or have systematic artifacts are not likely to be much improved by artifact rejection. (J. C. Hansen, unpublished software documentation)*

Hansen made this point in the context of artifact rejection, but it applies broadly to all post-processing procedures that are designed to clean up the data, ranging from averaging to filtering to independent components analysis. Some post-processing procedures are essential, but they cannot turn bad data into good data. You will always save time in the long run by eliminating electrical noise at the source, by encouraging subjects to minimize bioelectric artifacts, and by designing experiments to produce large effects.

### Active and Reference Electrodes

#### Voltage as a Potential Between Two Sites

Voltage can be thought of as the potential for current to move from one place to another, and as a result there is really no such thing as a voltage at a single point (if this fact is not obvious to you, you should read appendix 1 before continuing). Consider, for example, a typical twelve-volt automobile battery that has a positive terminal and a negative terminal. The voltage measurement of twelve volts represents the potential for current to move from the positive terminal to the negative terminal, and it doesn't make sense to talk about the voltage at one terminal in isolation. For example, you could touch one terminal without being shocked (assuming you weren't touching any other conductors), but if you touch both

terminals you will definitely receive a shock. Similarly, you can never record the voltage at a single scalp electrode. Rather, the EEG is always recorded as a potential for current to pass between two electrodes.

In household electrical systems, a metal stake driven deep into the ground beneath the house serves as an important reference point for electrical devices. The ground literally provides the reference point, and the term *ground* is now used metaphorically in electrical engineering to refer to a common reference point for all voltages in a system. I will therefore use the term *ground* in this general sense, and I will use the term *earth* to mean a stake driven into the ground.

If we measured the electrical potential between an electrode on a subject's scalp and a stake driven into the ground, the voltage would reflect any surplus of electrical charges that had built up in the subject (assuming the subject was not touching a conductor that was connected to earth), and this *static electricity* would obscure any neural signals. We could put an electrode somewhere on the subject's body that was connected to earth, and this would cause any static electricity in the subject to discharge into the earth, eliminating static differences and making it easier to measure changes in neural signals over time. However, it is dangerous to directly connect a subject to earth, because the subject might receive a dangerous shock if touched by an improperly grounded electrical device (such as a button box used for recording responses).

It is possible to create a virtual ground in the amplifier's circuitry that is isolated from earth and connect this ground to a ground electrode somewhere on the subject. You could then record the voltage between a scalp electrode and this ground electrode. However, voltages recorded in this way would still reflect electrical activity at both the scalp electrode and the ground electrode, so it would not provide some sort of absolute measure of electrical activity at the scalp electrode. Moreover, any environmental electrical noise that the amplifier's ground circuit picks up would in-

fluence the measured voltage, leading to a great deal of noise in the recording.

To solve the problem of the ground circuit picking up noise, EEG amplification systems use *differential amplifiers*. A differential amplifier uses three electrodes to record activity: an *active* electrode (A) placed at the desired site, a *reference* electrode (R) placed elsewhere on the scalp, and a ground electrode (G) placed at some convenient location on the subject's head or body. The differential amplifier then amplifies the difference between the AG voltage and the RG voltage (AG minus RG). Ambient electrical activity picked up by the amplifier's ground circuit will be the same for the AG and RG voltages and will therefore be eliminated by the subtraction.

It should now be clear that an ERP waveform does not just reflect the electrical properties at the active electrode but instead reflects the difference between the active and reference sites. There is simply no such thing as "the voltage at one electrode site" (because voltage is a potential for charges to move from one place to another). However, there is one way in which it would be meaningful (if slightly imprecise) to talk about voltages at individual sites. Specifically, if one could measure the potential for charges to move from a given point on the scalp to the average of the rest of the surface of the body, then this would be a reasonable way to talk about the voltage at a single electrode site. As I will discuss later in this chapter, points along the body beneath the neck do not matter very much in terms of electrical activity generated by the brain, so we could simplify this by referring to the voltage between one site and the average of the rest of the head. When I later speak of the absolute voltage at a site, I am referring to the potential between that site and the average of the entire head. It is important to keep in mind, however, that absolute voltages are just an idealization; in practice, one begins by recording the potential between two discrete electrode locations.


Originally, researchers assumed that the active electrode was near the active tissue, and the reference electrode was at some distant,

electrically neutral site. As the activity near the active electrode changed, researchers assumed that this would influence the voltage at the active site but not at the reference site. When obtaining recordings from several active sites, the same reference electrode is typically used for all of them. For example, you might place electrodes at active locations over frontal, parietal, and occipital cortex and use an earlobe electrode as the reference for all of them.

Figure 3.1 illustrates this, showing the distribution of voltage over the scalp for a single generator dipole (relative to the average of the entire head). The problem with this terminology is that there are no electrically neutral sites on the head (i.e., no sites that are equivalent to the average of the entire scalp), and the voltage recorded between an active site and a so-called reference site will reflect activity at both of the sites. In fact, one of the most important principles that I hope to convey in this chapter is that *an ERP waveform reflects the difference in activity between two sites, not the activity at a single site*. There are some exceptions, of course, which I will describe later in the chapter.

#### Choosing a Reference Site

Many researchers have tried to minimize activity from the reference electrode by using the most neutral possible reference site, such as the tip of the nose, the chin, the earlobes, the neck, or even the big toe. However, although some of these are useful reference sites, they do not really solve the problem of activity at the reference site. Consider, for example, the tip of the nose. This seems like it ought to be a fairly neutral site because it is fairly far away from the brain. To make the example more extreme, imagine a head with an extremely long nose, like that of Pinocchio. Pinocchio's nose is like a long wire attached to his head, and the voltage will therefore be approximately the same anywhere along this wirelike nose (see figure 3.1B). It doesn't really matter, therefore, whether you place the reference electrode at the tip of the nose or where the nose joins the head—the voltage is equal at both sites,



**Figure 3.1** Active and reference electrodes. (A) Example of an equivalent current dipole (arrow) inside a spherical head, with the resulting surface voltages (+ and - for positive and negative) on the surface. The recorded voltage will be the difference between the voltage at the active and reference electrodes. (B) Example of the use of a distant reference source. If the active electrode is at point A, it will not matter whether point B or point C is used as the reference, because the voltage at point C will be the same as at point B.

and so the difference between an active site and the reference site will be the same no matter where along the nose the reference electrode is. Because there is no reason to believe that the place where the nose joins the head is more neutral than any other part of the head, the tip of the nose also isn't more neutral than any other part of the head. This is not to say that the tip of the nose is an inappropriate site for a reference electrode. Rather, my point here is that there is no such thing as an electrically neutral reference site, so you must always keep in mind that an ERP waveform reflects contributions from both the active site and the reference site.

What, then, is the best site to use as a reference? There are three main factors involved in choosing a reference. First, given that no site is truly neutral, you might as well choose a site that is convenient and comfortable. The tip of the nose, for example, is a somewhat distracting place for an electrode. Second, you will want to avoid a reference site that is biased toward one hemisphere. For example, if you use the left earlobe as the reference, then the voltage recorded at left hemisphere electrodes will likely be different from the voltage recorded at the right hemisphere. You should avoid this sort of bias. Third, because an ERP waveform for a given active site will look different depending on the choice of the reference site, it is usually a good idea to use the same site for all of your experiments and to use a site that other investigators commonly use. This makes it easier to compare ERP waveforms across experiments and across laboratories.

The most common reference sites in cognitive neuroscience are the earlobes and the mastoid process (the bony protrusion behind each ear). They are close enough to each other that the resulting ERP waveforms should look about the same no matter which you use. The fact that they are commonly used is actually a good reason to continue using them, because this facilitates comparing data across experiments and laboratories. Thus, these sites satisfy the last of the three criteria described in the preceding paragraph. They are also convenient to apply and are not distracting, satisfy-

ing the first of the three criteria. In my lab, we use the mastoid rather than the earlobe because we find that an earclip electrode becomes uncomfortable after about an hour and because we find it easier to obtain a good electrical connection from the mastoid (because the skin is not so tough). However, I have heard that one gets a better connection from the earlobe in subjects with questionable personal hygiene. Thus, you should use whichever of these is most convenient for you.

Both the mastoid and the earlobe fail to satisfy the second criterion, because you must pick one side, leading to an imbalance between active electrodes over the left and right hemispheres. The simplest way to avoid this bias is to place electrodes at both the left mastoid (Lm) and the right mastoid (Rm) and then physically connect the wires from the Lm and Rm. This is called a *linked mastoids* reference, and it is not biased toward either hemisphere (you can do the same thing with earlobe electrodes). However, physically linking the wires from Lm and Rm creates a zero-resistance electrical bridge between the hemispheres, which distorts the distribution of voltage over the scalp and reduces any hemispheric asymmetries in the ERPs. To avoid these problems, it is possible to mathematically combine Lm with Rm by using an *average mastoids* reference derivation, which references the active site to the average of Lm and Rm (Nunez, 1981).

There are several ways to do this, and here I'll describe how we do it in my lab. When we record the EEG, we reference all of the scalp sites to Lm, and we also place an electrode on Rm, referenced again to Lm. After recording and averaging the data, we then compute the average mastoids derivation for a given site using the formula  $a' = a - (r/2)$ , where  $a'$  is the desired averaged mastoids waveform for site A,  $a$  is the original waveform for site A (with an Lm reference), and  $r$  is the original waveform for the Rm site (with an Lm reference).

Let me explain how this formula works. Remember that the voltage at a given electrode site is really the absolute voltage at the

active site minus the absolute voltage at the reference site. In the case of an active electrode at site A and an Lm reference, this is  $A - Lm$ . Similarly, the voltage recorded at Rm is really  $Rm - Lm$ . The average mastoids reference that we want to compute for site A is the voltage at site A minus the average of the Lm and Rm voltages, or  $A - ((Lm + Rm)/2)$ . To compute this from the data recorded at A and Rm, we just use some simple algebra:

$$a = A - Lm$$

Voltage recorded at site A is the absolute voltage at A minus the absolute voltage at Lm.

$$r = Rm - Lm$$

Voltage recorded at Rm is the absolute voltage at Rm minus the absolute voltage at Lm.

$$a' = A - (Lm + Rm)/2$$

Average reference derivation at site A is the absolute voltage at site A minus the average of the absolute voltages at Lm and Rm. This is just an algebraic reorganization of the preceding equation.

$$a' = A - Lm/2 - Rm/2$$

$$a' = A - (Lm - (Lm/2)) - (Rm/2)$$

This works because  $Lm/2 = Lm - Lm/2$ .

$$a' = (A - Lm) - ((Rm - Lm)/2)$$

This is just an algebraic reorganization of the preceding equation.

$$a' = a - (r/2)$$


Here we've substituted  $a$  for  $(A - Lm)$  and  $r$  for  $(Rm - Lm)$ .

In other words, you can compute the voltage corresponding to an average mastoids reference for a given site simply by subtracting half of the voltage recorded from the other mastoid (the same thing can be done with earlobe reference electrodes). All things considered, this is the best reference scheme for the majority of ERP experiments in the area of cognitive neuroscience.

### Alternatives to Traditional Reference Sites

There are two additional methods that researchers commonly use to deal with the problem of the reference site. First, imagine that you placed electrodes across the entire surface of the head. By using the average voltage across all of the electrodes as the reference, you could obtain the absolute voltage at each electrode site, and you wouldn't have to worry about the whole reference electrode issue. The mathematics of this would be trivial: the absolute voltage at a given site can be obtained by simply subtracting the average of all of the sites from each individual site, assuming that all sites were recorded with the same reference electrode. Although this would be ideal, it isn't practical for the simple reason that the neck and face get in the way of putting electrodes over about 40 percent of the head.

Some investigators use the average across all of the electrodes as the reference even though they don't have electrodes covering the entire head, and this can lead to some serious misinterpretations (see Desmedt, Chalklin, & Tomberg, 1990). Figure 3.2 illustrates this, showing illustrative waveforms recorded at the Fz, Cz, and Pz electrode sites. The left panel displays the absolute voltage at each site, and a positive-going wave can be seen in condition A relative to condition B at the Fz site. The middle panel shows the voltages that would be obtained using the average of the three electrodes as the reference. When the average of all sites is used as a reference, the average voltage across sites at any given time point is necessarily zero microvolts, so an increase in voltage at one site artificially induces a decrease in voltage at the other sites (the voltages across the entire head also sum to zero, but this is due to physics and not an artificial referencing procedure). As the figure shows, this referencing procedure seriously distorts the waveforms. For example, the short-duration, positive-going peak at around 400 ms at the Fz electrode site becomes a long-duration, negative-going peak when one uses the average reference. This occurs because of the large P3 wave at Pz in the absolute voltage waveforms; to achieve an



**Figure 3.2** Effects of using the average of all sites as the reference. The left column shows the absolute voltage recorded at each of three sites (Fz, Cz, and Pz), and the middle column shows the waveforms that are obtained when using the average of these three sites as the reference. The waveforms are highly distorted when using the average reference. The right column shows the results of using a somewhat larger set of electrodes, covering occipital and temporal electrodes as well as frontal, central, and parietal electrodes. Note that this is not a significant problem when sampling a large proportion of the head's surface (> 60 percent) with a dense set of electrodes (see Dien, 1998). Negative is plotted upward.

average voltage of zero, a large negative voltage has to be added onto the Fz site in the average reference waveforms.

The use of an average reference can lead to serious errors in interpreting ERP data. For example, I have reviewed several manuscripts in which an average-electrodes reference was used and the authors made a great deal out of the finding that an experimental effect was reversed in polarity at some sites relative to others. But this is necessarily the case when using the average across sites as the reference. Thus, it is very dangerous to use the average of all the electrodes as the reference. However, there may be conditions under which the average-electrodes reference might be appropriate. Specifically, Dien (1998) has argued that it is possible to obtain a close approximation of the true average of the entire head by using a large array of electrodes that covers most of the accessible part of the head, and the averaged-electrode reference may therefore be the best approach when enough electrodes are used (assuming that a consensus can be reached about exactly what set of electrodes is necessary to reach an adequate approximation of the true average voltage).

Another approach to this problem is not to rely on voltage measurements, but instead to examine the current flowing out of the head at each point, which does not depend on measuring a difference between an active site and a reference site (unlike voltage, current flow can be legitimately measured at a single point). Specifically, it is possible to convert the voltages measured in standard ERP recordings into current density (sometimes called scalp current density, SCD, or current source density, CSD). Current density is computed for a given electrode site on the basis of the distribution of voltage across the scalp. Technically speaking, one calculates the current density at the scalp by taking the second derivative of the distribution of voltage over the scalp (for a detailed description, see Pernier, Perrin, & Bertrand, 1988). Because ERPs are recorded at a finite set of discrete electrodes, the current density can only be estimated, but given a reasonably

dense set of electrodes, the estimate can be quite accurate. Moreover, it is not necessary to cover the entire head or even most of the head. If the electrodes are confined to a circumscribed region (e.g., the posterior 50 percent of the head), you can compute accurate estimates of current density within that region, although you shouldn't place much faith in the estimates at the outermost electrodes. Although current density has several advantages, it has one important disadvantage: it is insensitive to dipoles that are deep in the brain and preferentially emphasizes superficial dipoles (this is because the current from a deep source dissipates widely over the entire scalp). Consequently, current density provides a less complete picture of brain activity than traditional voltage measures. In addition, current density is an estimated quantity that is one step removed from the actual data, and it is therefore usually a good idea to examine both the voltage waveforms and the current density waveforms.

### Electrical Noise in the Environment

The voltage fluctuations of the scalp EEG are tiny (typically less than 1/100,000th of a volt), and the EEG must be amplified by a factor of 10,000–50,000 before it can be accurately measured. There are many sources of electrical activity in a typical laboratory that are much larger than the EEG, and these large electrical sources can produce small voltage fluctuations in the subject, in the electrodes, and in the wires leading from the subject to the amplifier. These induced voltage changes can be quite considerable when they are amplified along with the EEG (although once the EEG has been amplified, the effects of additional induced voltages are usually insignificant). Although some of the induced electrical noise can be eliminated by filters and other postprocessing techniques, it is always best to eliminate noise at the source (this is a corollary of Hansen's axiom). In this section, I will describe the major sources of electrical noise and discuss strategies for minimizing noise in ERP recordings.

An oscillating voltage in a conductor will induce a small oscillating voltage in nearby conductors, and this is how electrical noise in the environment shows up in the EEG. There are two major sources of oscillating voltages in a typical ERP lab, namely AC line current and video monitors. AC line current consists of sinusoidal oscillations at either 50 Hz or 60 Hz, depending on what part of the world you live in, and this can induce 50- or 60-Hz *line noise* oscillations in your EEG recordings. Video monitors may operate at a refresh rate of anywhere between 50 and 120 Hz (60–75 Hz is typical), but the resulting noise is often spiky rather than sinusoidal. Noise from the video monitor is especially problematic because the stimuli are time-locked to the video refresh, so the noise waveform is the same on every trial and is not reduced by the averaging process.

There are several things that you can do to decrease these sources of noise. The most common approach is to use the amplifier's filters to attenuate the noise. In most cognitive experiments, the ERPs of interest are composed mostly of frequencies under about 30 Hz, so you can filter out everything above 30 Hz (including line noise and video noise) without attenuating the ERPs very much. In addition, many amplifiers have a *line filter* that specifically filters out 50-Hz or 60-Hz noise. However, as chapter 5 will discuss, filters are a form of controlled distortion and their use should be minimized (line filters are especially bad), so filtering alone is not the best way to control noise.

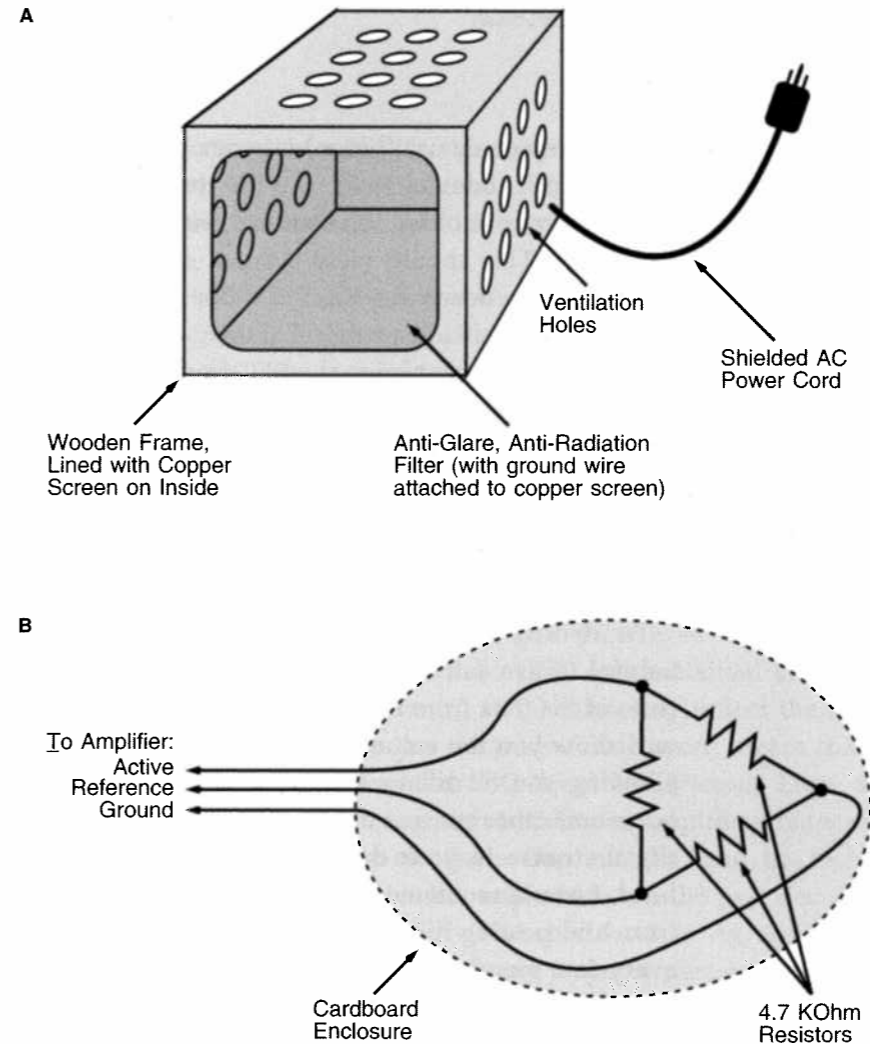
In many laboratories, the subject is placed in an electrically shielded chamber to minimize noise. This can be very effective, but only if there are no noise sources inside the chamber. For example, putting the video monitor inside the chamber creates so much electrical noise that it is hardly worth having a shielded chamber (except in environments with other very large sources of noise). Two approaches are commonly used for solving this problem. First, you can place the video monitor just outside a window in the chamber (but the window must be made of specially treated shielded glass). Second, you can place the monitor inside a



*Faraday cage* inside of the recording chamber. Figure 3.3A shows a design for a Faraday cage that can be built very easily for a relatively small amount of money (US\$200–400). This cage consists of copper screen shielding surrounded by a wooden exterior (with ventilation holes). A shielded piece of glass is placed at the front so that the subject can see the front of the monitor. Shielded glass is available from computer accessory suppliers. There are several different types, and you should get one that has a ground wire coming out of it that you can attach to the copper shielding. You should also get a shielded AC power cord, available at electronics supply companies. A well-shielded Faraday cage can dramatically reduce electrical noise in the EEG, and it is well worth the modest expense.

There may also be other sources of electrical noise inside the chamber. For example, AC lights can be a problem, so you can replace them with DC lights of the type found inside an automobile or recreational vehicle. Basically, you want to make sure there is nothing inside the chamber that might create electrical noise, especially devices powered by AC line voltage. Some of these are obvious, such as lights, but others are not. For example, when I was putting together my current laboratory, I was surprised to find that the cables leading from a stereo amplifier outside the chamber to speakers inside the chamber created significant electrical noise in the EEG—encasing these wires inside shielded conduits eliminated the noise.

Fortunately, there is a fairly easy way to track down sources of electrical noise such as this. As figure 3.3B illustrates, you can create a simulated head out of three resistors, and by connecting this to your amplifier and digitization system, you can see how much noise is present at different locations inside the recording chamber or room. First, place the simulated head where the subject sits and connect the fake head to your amplifier and digitization system using a very high gain (50,000 or 100,000), a fairly high sampling rate (e.g., 1000 Hz), and wide open filter settings (e.g., 0.01–300 Hz). You should see clear signs of electrical noise (including a sinusoi-



**Figure 3.3** (A) Faraday cage that can be used to reduce electrical noise from a video monitor. It consists of a plywood box with ventilation holes, completely lined with copper screen (copper sheets may be used on surfaces without ventilation holes). The front of the cage has a glare filter that is covered with a conductive film, which is connected to the copper shielding. A power strip is placed inside the cage (not shown); this is connected to a shielded AC power cord. (B) Simple simulated head for use in finding noise sources. It consists of three 4.7 KOhm resistors (or similar values), connected to form a triangle. Each corner of the triangle is connected to wires leading to the active, reference, and ground inputs of the EEG amplifier. The set of resistors should be enclosed in a nonconductive substance (such as cardboard) for ease of handling.

dal line-frequency oscillation) on your EEG display. Then, turn off every electrical device in or near the recording chamber, including any AC power entering the chamber, such as the power for the fan and lights. This should yield a much cleaner signal on your EEG display (i.e., a nearly flat line). If it doesn't, then you may be picking up noise outside of the chamber, in which case you should shield the cables that bring the EEG from the chamber to your digitization system. Once you have obtained a clean signal, you can turn the electrical devices back on one at a time, noting which ones create noticeable noise. You may also find it useful to place the simulated head near wires and devices that are potential noise sources, which will make it easier to see any noise that they generate.

To identify and eliminate sources of noise in this manner, it is helpful to use software that computes an on-line frequency spectrum of the data from the simulated head. The frequency spectrum will show you the amount of activity in various frequency bands, allowing you to differentiate more easily between line-frequency noise and other sources of noise. This may also allow you to quantify the noise in your data. Once you have minimized the noise level, I would recommend making a printout of the frequency spectrum and posting it in the lab. You can then check the noise level every few months and compare it with this printout to see if the noise level has started to increase (e.g., because the video monitor has been replaced with one that generates more noise).

### Electrodes and Impedance

Now that we have discussed the nature of the voltages that are present at the scalp, let's discuss the electrodes that pick up the voltages and deliver them to an amplifier. Basically, a scalp electrode is just a conductive material attached to a wire. In most cases, an electrode is a metal disk that forms an electrical connection with the scalp via a conductive gel. The choice of metal is fairly important, because some metals corrode quickly and loose


their conductance. In addition, the circuit formed by the skin, the electrode gel, and the electrode can function as a capacitor that attenuates the transmission of low frequencies (i.e., slow voltage changes).

Until the 1980s, most researchers used silver electrodes covered with a thin coating of silver-chloride (these are typically called *Ag/AgCl* electrodes). These electrodes have many nice properties, but they can be somewhat difficult to maintain. In the 1980s, many investigators started using electrode caps made by Electro-Cap International, which feature tin electrodes that are very easy to maintain. In theory, tin electrodes will tend to attenuate low frequencies more than *Ag/AgCl* electrodes (Picton, Lins, & Scherg, 1995), but Polich and Lawson found essentially no difference between these two electrode types when common ERP paradigms were tested, even for slow potentials such as the CNV and sustained changes in eye position (Polich & Lawson, 1985). This may reflect the fact that the filtering caused by the electrodes is no more severe than the typical filter settings of an EEG amplifier (Picton, Lins, & Scherg, 1995). Moreover, using high-impedance amplifiers reduces the filtering properties of an electrode. On the other hand, the technology for *Ag/AgCl* electrodes has improved over the past decade, and now many investigators have switched back to *Ag/AgCl*. Either tin or *Ag/AgCl* should be adequate for most purposes, unless you are recording DC potentials.

Because electricity tends to follow the path of least resistance, it is important to ensure that the resistance between the electrode and the scalp is low. Technically speaking, the term *resistance* applies only to an impediment to direct current (DC), in which the voltage does not change over time. When the voltage varies over time (i.e., alternating current or AC), the current can be impeded by *inductance* and *capacitance* as well as resistance; the overall impediment to current flow is called *impedance* (see the appendix for more details). Thus, it is more proper to use the term impedance rather than resistance in the context of ERP recordings, in which the voltage fluctuates over time. Impedance is frequently

**Box 3.1** Electrode Naming and Placement Conventions

This section will describe the most common system for placing and naming electrode sites, which was developed in the late 1950s by the International Federation of Clinical Neurophysiology (Jasper, 1958). This system is called the *10/20 system*, because it places electrodes at 10 percent and 20 percent points along lines of latitude and longitude, as illustrated in the figure. The first step in this system is to define an equator, which passes through the nasion (the depression between the eyes at the top of the nose, labeled Nz), the inion (the bump at the back of the head, labeled Iz), and the left and right pre-auricular points (depressions just anterior to the middle of the pinnae, labeled A1 and A2). A longitude line is then drawn between Iz and Nz, and this

**Box 3.1** (continued)

line is then divided into equal sections that are each 10 percent of the length of the line. Additional latitude lines, concentric with the equator, are then placed at these 10 percent points. Most of the electrode sites can then be defined as points that are some multiple of 10 percent or 20 percent along these latitude lines. For example, the F7 electrode is on the left side of the latitude line that is 10 percent from the equator, 30 percent of the distance around the latitude line from the middle. The exceptions to this 10/20 rule are F3 and F4 (halfway between Fz and F7 or F8) and P3 and P4 (halfway between Pz and T5 or T6).

Each electrode name begins with one or two letters to indicate the general region of the electrode (Fp = frontal pole; F = frontal; C = central; P = parietal; O = occipital; and T = temporal). Each electrode name ends with a number or letter indicating distance from the midline, with odd numbers in the left hemisphere and even numbers in the right hemisphere. Larger numbers indicate greater distances from the midline, with locations on the midline labeled with a "z" for zero (because the number 0 looks too much like the letter O).

This scheme has been extended to include many more additional electrode sites (American Encephalographic Society, 1994; Klem et al., 1999), and many individual investigators have developed their own naming schemes. You can use any system you like, as long as it accommodates differences in head sizes and is described in a manner that can be related to the 10/20 system.

denoted by the letter Z and measured in units of Ohms ( $\Omega$ ) or thousands of Ohms ( $K\Omega$ ) (e.g., "the Z for each electrode was less than 5  $K\Omega$ ").

It is common practice in ERP research to reduce the impedance of the skin to below 5  $K\Omega$  before attaching the electrodes; in fact, this practice is so common that I usually neglect to mention it when I write journal articles (although it is definitely worth mentioning). To reduce the impedance, it is necessary to remove the outer layer of dead skin cells that are primarily responsible for the naturally high impedance of the skin. There are two main classes of methods for accomplishing this, and the choice of method depends on the type of electrode being used. If the electrodes are

attached to the head individually using some sort of adhesive, you can first clean the skin at each site with an alcohol pad and then rub it with an abrasive paste (it is also possible to do this in one step with alcohol pads that include an abrasive, which I recommend). If you are using an electrode cap, it is not usually possible to abrade the electrode sites before applying the electrodes. Instead, you must insert an abrading implement through the hole in each electrode to abrade the underlying skin. Most laboratories use a blunt needle or the wooden end of a cotton-tipped swab for this. In my lab, we use a sharp, sterile needle. This sounds painful, but it is actually the least painful technique when done properly. You should not stick the needle into the skin, but instead rub it gently along the surface of the skin to displace the top layer of dead skin cells. When done properly, the subject can barely feel the needle. Of course, it is important to thoroughly disinfect the electrodes after each subject, but this is true no matter how one abrades the skin.

Although most investigators know that reducing the impedance is important for obtaining a clean signal, many do not know the precise nature of the problems that high impedance can create. There are really two main types of problems that high impedance can create: (1) decreased common-mode rejection, and (2) increased skin potentials.

First, let's consider common mode rejection. As we discussed near the beginning of this chapter, EEG amplification is accomplished by means of differential amplifiers that amplify the difference between the active-ground voltage and the reference-ground voltage. This subtracts away any electrical noise that is present in the ground (as well as any noise that is equal in the active and reference electrodes) and is essential for obtaining clean EEG recordings. Unfortunately, it is not trivial to produce an amplifier that performs this subtraction perfectly. If one of the two signals in the subtraction is attenuated slightly, then the subtraction does not work and the noise will not be completely subtracted away. The ability of an amplifier to subtract away environmental noise accu-

rately is called *common mode rejection*, and it is usually measured in decibels (dB; an exponential scale in which a doubling of power equals an increase of 3 dB). A good EEG amplifier will have a common mode rejection of at least 70 dB, and preferably over 100 dB. For reasons too complex to explain here, common mode rejection becomes less effective when the impedance is higher, and it can become extremely bad when the impedance is generally high and differs considerably among the active, reference, and ground electrodes. Thus, low electrode impedance helps you avoid picking up environmental noise.

Common mode rejection depends not only on the impedance of the electrodes, but also on the input impedance of the amplifier. An amplifier with a very high input impedance can tolerate higher impedance electrodes while maintaining a good level of common mode rejection. However, high amplifier impedance cannot solve the second problem associated with high electrode impedance, namely *skin potentials*. There is a tonic electrical potential between the surface of the skin and the deep layers of the skin, and this voltage changes whenever the skin's impedance changes. For example, when a subject sweats, the skin's impedance changes and the voltage at the surface changes. Similarly, if the subject moves, shifting the electrode to a slightly different position with a slightly different impedance, the voltage will change. These changes in voltage are called skin potentials, and because they are often very large, they can be a major source of low-frequency noise in ERP recordings. Picton and Hillyard (1972) showed that decreasing the impedance of the skin dramatically reduces skin potentials, and this provides a second compelling reason for decreasing the impedance before recording ERPs.

As computers and electronics have become cheaper and more powerful, ERP researchers have used more and more electrodes in their recordings. In the 1970s, most laboratories used one to three active scalp electrodes, but most laboratories now have the capability to record from at least thirty-two electrodes and many can record from 128 or even 256 electrodes.

Some special problems arise when using dense arrays of electrodes, and I will describe the most important ones here. The biggest drawback to large arrays of electrodes is the time required to decrease the impedance at each site. If it takes an average of one minute to decrease the impedance of each electrode, then a 64-channel electrode cap will require over an hour. A second problem is that the electrode gel inside a given electrode may leak out and cause an electrical bridge with nearby electrodes, distorting the scalp distribution of the electrical potentials. Because the whole point of recording from a dense electrode array is to measure the scalp distribution more accurately, this is a significant problem. The third problem is simply that as the number of electrodes increases, the probability of a bad connection increases and the probability of the experimenter noticing a bad connection decreases.

The geodesic sensor net Tucker and his colleagues developed provides a means of recording from many sites while minimizing some of these problems (Tucker, 1993). The electrodes in this device are basically just sponge-tipped tubes filled with saline, and the electrodes are connected to special high-impedance amplifiers so that no abrasion of the skin is necessary to obtain a reasonable level of common-mode rejection. Consequently, it is possible to attach 128 electrodes in a matter of minutes. Fast application, high-impedance recording systems of various types are now available from several companies. It might appear that these systems solve the problem of the long amount of time required to abrade a large number of sites (and they clearly reduce the risk of disease transmission through the electrodes). However, as discussed above, skin potentials are much larger in high-impedance recordings, and this can be a very considerable source of noise. Moreover, these systems do not solve the problem of a greater likelihood of a bad connection with a larger number of electrodes and the problem of bridging across electrode sites (although some systems can automatically detect bridging).

These high-impedance systems trade speed of application for a degradation in signal quality. Is this a worthwhile tradeoff? The

answer to this question depends on how much the signal is degraded and how important it is to record from a large number of electrodes (and to apply them quickly). The issue of signal degradation in these systems has not, to my knowledge, been systematically explored in papers published by independent sources, although there are some publications by people associated with the companies that sell these systems. For example, Tucker's group performed a study of the effects of electrode impedance on noise levels (Ferree et al., 2001), and they found only a modest and statistically insignificant increase in line-frequency (60-Hz) noise as impedance increased. However, they are not unbiased, so they may not have sought out the conditions most likely to produce high noise levels. Moreover, they eliminated skin potentials by filtering out the low frequencies in the signal, making it impossible to assess the effects of high impedance on this very significant source of noise. They argued that 60-Hz noise and skin potentials are not a problem given modern methods, but that has certainly not been my experience!

When making decisions about EEG recording systems, you should keep in mind two important factors. First, as I will discuss in chapter 4, you can offset a decrease in the signal-to-noise ratio by including more trials in your averages, but a very large number of additional trials will be necessary to offset even a modest decrease in the signal-to-noise ratio. For example, you would need to double the number of trials to offset a 30 percent decrease in the signal-to-noise ratio. Thus, you might be able to save an hour by using an easy-to-apply electrode system, but you might need to spend an additional two or three hours collecting data to offset the resulting reduction in signal-to-noise ratio. Before using a fast procedure for applying electrodes, you must be very careful to evaluate the hidden costs that these procedures may entail.

A second key factor is that large numbers of electrodes (> 40) are only occasionally useful in ERP experiments. As discussed in chapters 1 and 7, it is extremely difficult to localize the neural generators of an ERP component solely on the basis of the distribution

**Box 3.2** My Personal Perspective on Large Electrode Arrays

I have seen many papers in the last few years in which ERPs were recorded from sixty-four or more electrodes even though the main conclusions of the paper did not necessitate large numbers of electrodes. And in most cases, the data were not very clean. Worse yet, the authors frequently included plots of the waveforms from each site, but given the large number of sites, each waveform was tiny and it was difficult to discern the details of the experimental effects. Because ERPs cannot typically be localized, recording from a large number of channels is usually more of a hindrance than a help. Personally, I would rather have clean data from ten electrodes than noisy data from a thousand electrodes. This brings us again to Hansen's axiom: There is no substitute for good data (and noisy data aren't improved by having lots of channels).

of voltage over the head, so large numbers of electrodes are useful primarily when you will be obtaining some form of converging evidence (e.g., MEG recordings or fMRI scans). Moreover, in the few situations in which large numbers of electrodes are useful, it is also imperative to obtain extremely clean data, which means that common mode rejection should be maximized and skin potentials should be minimized. Thus, there are large costs associated with large electrode arrays, and one should use them only when the benefits clearly outweigh the costs.

**Amplifying, Filtering, and Digitizing the Signal**

Once the electrodes have picked up the EEG, it must be amplified and then converted from a continuous, analog voltage into a discrete, digital form that a computer can store. Fortunately, these processes are relatively straightforward, although there are a few important issues to consider, such as selecting an amplifier gain and choosing a digitization rate. Although the EEG is amplified before it is digitized, I will discuss the digitization process first because the settings you will use for your amplifier will make more sense once we have discussed the digitization process.

**Analog-to-Digital Conversion and High-Pass Filters**

A device called an *analog-to-digital converter* (ADC) converts EEG voltage fluctuations into numerical representations. In most EEG digitization systems, the ADC has a resolution of twelve bits. This means that the ADC can code  $2^{12}$  or 4096 different voltage values (intermediate values are simply rounded to the nearest whole number). For example, if the ADC has a range of  $\pm 5$  V, a voltage of  $-5$  V would be coded as 0, a voltage of  $+5$  V would be coded as 4096, and the intermediate voltages would be coded as  $4096 \times ((V + 5)/10)$ , where  $V$  is the voltage level being digitized. Typically, voltages that fall outside the range of the ADC will be coded as 0 for negative values and 4096 for positive values, and you'll obviously want to avoid exceeding this range. Generally, you will want to set the gain on your amplifier so that the range of the ADC is rarely or never exceeded (and, as described in chapter 4, you will want to discard trials that exceed the ADC range). The fact that the EEG is digitized with only twelve bits of resolution does not mean that your data will ultimately be limited to twelve bits of resolution. When you average together many trials, the resolution increases greatly, so twelve bits is sufficient for most cases.

There are two settings on a typical EEG amplifier that will affect whether or not you will exceed the range of the ADC (or the range of the amplifier itself). The first factor is the gain: as you increase the gain, you increase the chance that the EEG will exceed the ADC's range (you don't want to set the gain too low, however, because this will cause a loss of resolution). The second factor is the setting of the high-pass filter. A high-pass filter is a device that attenuates low frequencies and passes high frequencies. High-pass filters are important, because they attenuate the effects of large gradual shifts in voltage due to skin potentials. Even if you use low-impedance recordings, some skin potentials will occur, and these potentials can cause the EEG to drift out of the range of the ADC and amplifier, causing a "flat-line" signal. The higher you set the frequency of the high-pass filter, the less drift will occur. However, as I will discuss further in chapter 5, filters always lead to

distortion of the ERPs, and the distortion produced by a high-pass filter almost always becomes worse as you increase the cutoff frequency of the filter. In the 1960s and early 1970s, for example, most ERP experiments used a cutoff frequency of 0.1 Hz, but researchers eventually realized that this cutoff frequency led to a significant reduction in the apparent amplitude of the P3 wave (Duncan-Johnson & Donchin, 1979). As a result, most investigators now use a cutoff of 0.01 Hz, and this is what I would recommend for most experiments. There are, however, some cases in which a higher cutoff frequency would be appropriate. For example, when using children or psychiatric/neurological patients as subjects, voltage drifts are very common, and eliminating these artifacts may be worth some distortion of the ERP waveforms.

In some cases, it is also desirable to record the EEG without using a high-pass filter (these are called direct coupled or DC recordings). Some digitization systems use sixteen-bit ADCs to avoid the problems associated with the lack of a high-pass filter. By using a sixteen-bit ADC, it is possible to decrease the gain by a factor of sixteen, which decreases the likelihood that the EEG will drift out of the range of the ADC. This reduction in gain is equivalent to losing four ADC bits, but this is compensated for by the additional four bits of resolution in the sixteen-bit ADC, yielding the equivalent of twelve bits of resolution. It is also possible to apply a digital high-pass filter to the data at the end of the experiment, which would give you access to both the original unfiltered data and filtered data. However, although this sounds like it would be the best of both worlds, it is difficult in practice to limit a digital filter to very low frequencies. Consequently, it is almost always best to use an amplifier with analog filters that can filter the very low frequencies, reserving DC recordings for the relatively rare cases in which DC data are important for reaching a specific scientific conclusion. I would definitely be wary of buying an amplification system that does not include analog high-pass filters; I have heard several ERP researchers express regret after purchasing DC-only amplifiers.

### Discrete EEG Sampling and Low-Pass Filters

The EEG is converted into a voltage at a sequence of discrete time points called samples. The *sampling period* is the amount of time between consecutive samples (e.g., 5 ms) and the *sampling rate* is the number of samples taken per second (e.g., 200 Hz). When many channels are sampled, they are scanned sequentially rather than simultaneously, but the digitization process is so fast that you can think of the channels as being sampled at the same time (unless you are examining extremely high-frequency components, such as brainstem evoked responses). How do you decide what sampling rate to use? To decide, you need to use the *Nyquist theorem*, which states that all of the information in an analog signal such as the EEG can be captured digitally as long as the sampling rate is at least twice as great as the highest frequency in the signal. This theorem also states that you will not only lose information at lower sampling rates, but you will also induce artifactual low frequencies in the digitized data (this is called *aliasing*).

To use the Nyquist theorem, you need to know the frequency content of the signal that you are recording so that you can set your sampling rate to be at least twice as great as the highest frequency in the signal. However, the raw EEG signal may contain noise at arbitrarily high frequencies, so you cannot digitize the raw EEG without risking aliasing. EEG amplifiers therefore contain low-pass filters that attenuate high frequencies and pass low frequencies, and your sampling rate will depend primarily on the cutoff frequency that you select for your low-pass filters. For example, many investigators set the cutoff frequency at 100 Hz and digitize at 200 Hz or faster. As I will discuss in chapter 5, a cutoff frequency of 100 Hz does not completely suppress everything above 100 Hz, so you should use a digitization rate of at least three times the cutoff frequency of the filter. In my laboratory, for example, we usually filter at 80 Hz and digitize at 250 Hz.

The Nyquist theorem gives us a precise means of determining the sampling rate given the filter's cutoff frequency, but how do you decide on the cutoff frequency? Unfortunately, the answer

to this question is not so straightforward because, as mentioned above, filters are guaranteed to distort ERP waveforms. In general, the higher the frequency of a low-pass filter, the less distortion it will create. Consequently, you will want to choose a fairly high cutoff frequency for your low-pass filter. However, you don't want your filter frequency to be too high, because this will require a very high digitization rate, leading to huge data files. The best compromise for most cognitive neuroscience experiments is a low-pass cutoff frequency between 30 and 100 Hz and a sampling rate between 100 and 300 Hz. If you are looking at the early sensory responses, you'll want to be at the high end of this range (and even higher if you wish to look at very high-frequency components, such as the brainstem evoked responses). If you are looking only at lower frequency, longer latency components, such as P3 and N400, you can set your cutoff frequency and sampling rate at the low end of this range. Keep in mind, however, that the lower your filter frequency and sampling rate, the less temporal precision you will have.

#### **Amplifier Gain and Calibration**

The signal from each electrode is amplified by a separate EEG channel. The gain (amplification factor) that you will use depends on the input range of your analog-to-digital converter. If, for example, your analog-to-digital converter allows an input range of  $-5\text{ V}$  to  $+5\text{ V}$ , you will want to set the gain of your amplifier so that its output is near  $-5\text{ V}$  or  $+5\text{ V}$  when it has the most extreme possible input values. That is, you will want to use the entire range of the analog-to-digital converter, or else you will not be taking advantage of its full resolution. Most systems work best with a gain somewhere between 1,000 and 50,000 (my lab uses 20,000).

Even if you select the same gain setting for all of the channels of your amplifier, the gains will probably not be exactly the same. It is therefore necessary to calibrate your system. The best way to do this is to pass a voltage of a known size through the system and

measure the system's output. For example, if you create a series of  $10\text{ }\mu\text{V}$  voltage pulses and run them into your recording system, it may tell you that you have a signal of  $9.8\text{ }\mu\text{V}$  on one channel and  $10.1\text{ }\mu\text{V}$  on another channel. You can then generate a scaling factor for each channel (computed by dividing the actual value by the measured value), and multiply all of your data by this scaling factor. You can do this multiplication on the EEG data or on the averaged ERP waveforms; the result will be the same.