With Photoshop, you can either work creatively/holistically with images, or mathematically/logically. You can use either side of your brain and come up with really fascinating results. Sometimes the results will be great, and just what you want, other times you'll get stuff meant for the digital garbage dump.

Photoshop tools have become a standard across a lot of image editing and drawing applications. And notice how similar they look to the old MacPaint tools we already saw.

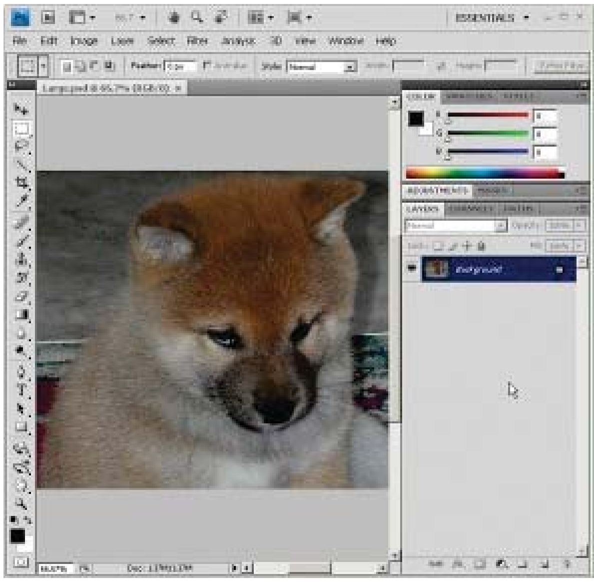

When learning the basics of Photoshop, try to translate from something you already know - probably a word processor is a good analogy. If you think in general terms, rather than following some online step-by-step tutorials, you might just find that you can learn a lot more than just Photoshop! This, of course, is the goal. Since you'll never be able to know the latest version of each and every piece of software, it's key to understand the fundamentals, and be able to translate this knowledge to new situations when needed.

When you later decide to try out video and audio editing software you'll find the tools will be somewhat different, but many have the same principles behind them. I'll try to point out the similarities as we go along. You'll discover many of your own, too.

And, Photoshop for the Mac and the PC work practically the same. So you'll have no problems switching back and forth between these two platforms. This is another reason why Photoshop is such a popular choice for graphics editing.

A few interesting resources

Here's an interesting look at Google streetview (which gives you a close up view of the world, including anything that is happening on the street when the little Google driver comes by documenting)
http://www.artfagcity.com/2009/08/12/img-mgmt-the-nine-eyes-of-google-street-view/

An a very neat to watch video timelapse that walks you through the creation of a Macworld magazine cover - from product photography to desktop publishing.
http://peterbelanger.com/posts/36-cover-creation

Audio Overview

Or how sound works in the physical (analog) world.

Waveforms
Every sound you hear is made up of sound waves traveling through the air, hitting your eardrums. The hairs in your eardrum vibrate according to this change in atmospheric pressure, and your brain interprets this as different sounds.

A sound wave (sine) looks like this:

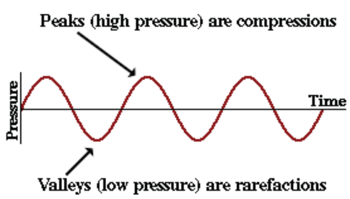

And acts like this on your eardrum:

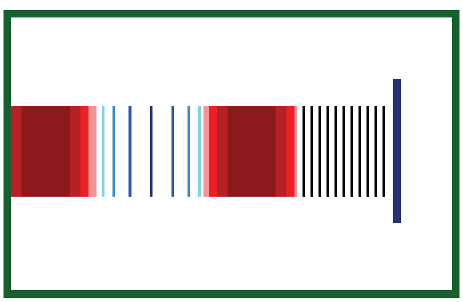

Sound acts the same way when it comes through your amplified speaker system. Take off the cover plate of your speaker, and watch the cone of your speaker move in and out, just like the above graphic. (You might have to turn up the volume a bit to really see it.)

The simplest of all sounds is that sine wave you see above. Yep, the same sine of sine, cosine, and tangent -- made famous from your math class. Luckily, you won't need to break out your calculator for this. Sine waves are based on circles. In audio, sine waves are the result of tracing a circle's circumference over time.

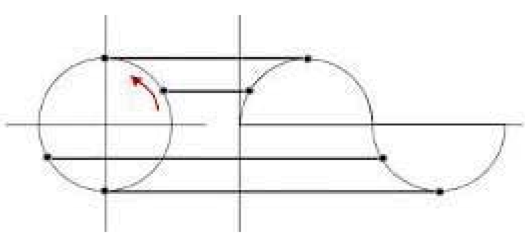

Can you see how the point on the circle goes around and, if drawn over time, creates a sine wave?
This sound wave, like any other wave, needs a medium to travel through. In most cases this is air, but, it

could be water or something else. Have you ever been in a pool and tried to talk to a friend underwater? It sounds weird, doesn't it? It sounds weird because sound waves travel through air at a faster rate than they do through water.

Frequency

Sound pitch is measured as frequency, or how many times per second the pattern of waves repeats. This is easiest explained with simple waves. Play the QuickTime movie below to see and hear what a sine wave sounds like as you double its frequency.
Notice that the higher the pitch is, the more sine wave "periods"of the wave you see?

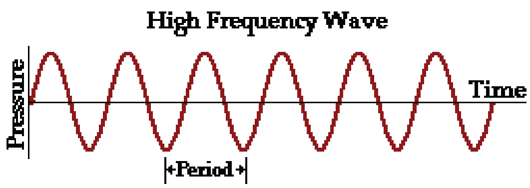

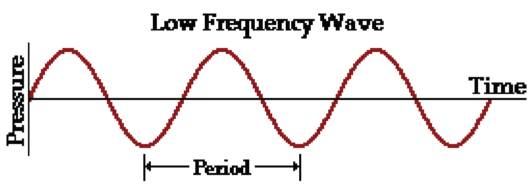

Hertz (Hz) is the measure of how many vibrations per second occur. It's often seen as KHz or **kilohertz** because, in audio, typically the waves are moving fast. For example, 20Hz is 20 vibrations per second, while 18.5KHz is 18,500 vibrations per second. A newborn baby, with no hearing damage, can hear just over 20KHz, while people with some hearing loss (due to normal living conditions), usually hear up to around 17KHz.

Typically, men lose some high-end hearing earlier than women. Pay attention here. What's the reason for so much of everyone's hearing loss? Headphones playing music way too loud right next to your eardrums. Especially the "ear-bud" style headphones that you jam right into your ear. Why? Remember those pressure waves that cause your eardrum to vibrate causing you to hear sound? Well, when you've got headphones in your ears, they can act like plungers, pushing lots of high pressure (peaks) into your eardrum, but preventing the equivalent low pressure (valleys) out of your ear. Just be careful those of you who listen to many hours of music on your portable players!

Timbre

Timbre (pronounced like tam-bur) is the tone color of a sound. So, not every tone with the same pitch sounds the same. Obvious, right? If you hear an orchestra tuning, you hear them all playing 440Hz (concert pitch A), but the oboe sounds different than the violins. That's because they have different "overtones" that go into their waveforms.

Amplitude

The amplitude, or height of a wave's peaks and valleys, determines how loud it sounds. This is referred to as dB or **decibels** in the audio world.
The human ear is incredibly dynamic in that we can hear a very large range of amplitudes. The faintest sound that you can just barely hear is called the "threshold of hearing". On the other end of the spectrum is the "threshold of pain". The most intense sound that the ear can safely detect without suffering physical damage is more than one billion times more intense than the quietest sound. That's a pretty dramatic range.

Because of this huge range, decibels are based on a logarithmic (or exponential) scale. The threshold of hearing is 0dB. A sound ten times more intense (like rustling leaves) is 10dB. Normal conversation is probably close to 60dB. A rock concert (by law) shouldn't ever exceed 115dB. The threshold of pain is about 130dB or 10^{13} times louder than the threshold of hearing.

Many years ago I used to do "sound control" at the auto races. We'd set up a microphone a specific distance from the track and make sure that no car went over 90 dB. If they did, that car got a technical black flag and a mandatory pit stop for them to fix the noise! The microphone was connected to an "SPL" or sound pressure level meter. This way we had a specific way to measure this, rather than just a subject feeling about it.

Resources
If you want to check out more about the physics of sound, which is fascinating and can help you make connections between math, computers, and physics (making an excellent multimedia generalist) try this site:
http://www.glenbrook.k12.il.us/gbssci/phys/Class/sound/soundtoc.html

Audio Overview
Or how sound works in the computer (digital) world.

Microphones and Speakers

Microphones and speakers have exactly opposite functions.
Microphones are designed to record an incoming audio signal (sound wave). They then convert the sound wave by use of a transducer into a corresponding electrical signal. Magnetic tape, for example, stores this electrical signal in "magnet" form. Your hard drive stores this signal as a string of binary numbers.
A speaker is designed to take an electrical signal (or wave) and convert it, through the very same transducer, into a corresponding audio sound wave that our ears hear as a "reproduction" of the original sound.

A **transducer** is at the heart of making amplified music. It converts one type of energy into another. The more accurately your microphone and speakers reproduce sound, the better your system and the closer your recording will sound to "real life". This sounds easy enough, but there are a lot of places along the signal flow that degradation of your audio signal can occur. Our job is to try to eliminate as much of this degradation as possible.

Mixers/Multitracks
Mixers, at their core, are designed to combine various audio signals and "mix" them down (typically) to a stereo output. For our sake, it's important to know about combining signals.

Ultimately, if you're sending a **stereo** signal through an audio device to someone's headphones or speakers, it means you can send them a Right and a Left channel of audio. Sometimes these two signals are slightly different from each other, and you can get a spatial sense, or placement of sound sources. For instance, when you listen to the radio do the vocals sound like they are coming from only one of your speakers? Or does it sound like they are coming from someplace in-between both of your speakers? It's the illusion of stereo to be able to place sounds (or pan them) from stereo right, to stereo left. A typical use of the stereo field is when you are doing a spoken-word interview with two people to have one person's voice coming from "stereo right" and the other person coming from "stereo left" to help the listener keep track of who is speaking.

But, what do you do when you are in a more complex recording environment. For example, maybe there are two mics recording the vocals, and three guitars that each had their own mic. For this you use mixers. Mixers let you to have 8, 12, 24 or more input channels, and ultimately, mix down to two outputs (stereo). This is because most people only have the capacity to play back two channels of audio (like on your home stereo, on your iPod, or your computer). Some systems allow for more than two channels - maybe your TV has a "surround sound" system? With Dolby 5.1 you get 5 channels of sound.

You should consider your audience when deciding how many channels of output you need. A mono audio file is half the file size of a stereo file. Let's find out why.

What is Digital Audio?

By encoding analog sound into digital format, you can edit and process sound digitally - using the exact same principles you use with your word processor. You can add effects and edit as many "generations" as you want, without any loss of quality. Just like you can have a text document, and make a copy of it, add bold and other formatting, cut out a few paragraphs, and have no loss in your text documents "quality", the same is true for your audio, if you know the right file formats to use.

Digital means Sampling

Unlike sampling color, where you simply analyze one particular pixel and determine what color you have, with audio you need to add the dimension of time.

A **sample** is a digital representation of the frequency of a sound at a particular moment in time. To represent a frequency you need to sample it at least two times a cycle. At its simplest, the sampling theorem states that in order to uniquely represent a continuous-time signal by a discrete set of samples, the natural frequency at which the signal is sampled must be greater than *twice* the highest natural frequency in the original signal. The minimum natural sampling frequency is often referred to as the **Nyquist** sampling rate.

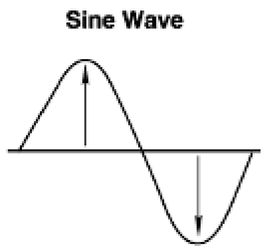

The math of this isn't as complex as you might think. If you are trying to represent the entire audio spectrum (20Hz-20KHz), according to the Nyquist theory, you need to sample at minimum 40,000 times per second (two times the highest frequency). The number of times per second you digitally "sample" a wave is called the **sampling rate**. The CD sampling rate slightly higher than this minimum, it is 44.1KHz. This means that samples are being taken 44,100 times per second to produce CD-quality audio.

If you sample a wave less then twice in one period, you get an aliasing effect that will not accurately represent the audio you are digitizing.

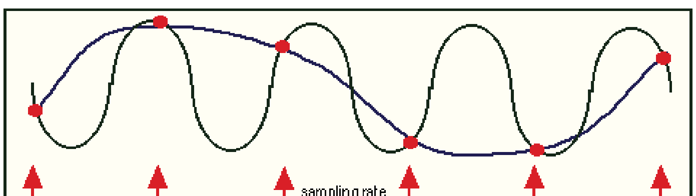

Aliasing produces waveform (blue) at 1/4 the original signal's frequency (green).

Want to hear how this sounds. Go to the below link and listen to a frequency sweep. It should just sound like it keeps getting higher and higher pitched, but doesn't. It's a slide from 440Hz up to 44,000 Hz (well above your hearing range). (Your doggie can hear frequencies this high!)
http://www.indiana.edu/~emusic/foldover.aif

For a visual analogy: In old movies, wheels on wagons sometimes appear to spin backwards. This is because the cameras did not sample the visual images fast enough to properly reconstruct them later during playback. Aliasing explains the reasons for this odd visual effects. In the graphic above see how the blue waveform is "aliased" so it appears to be going much slower than the green wave.

Bit depth is simply the "depth" of sound. Either 8bit, 16bit, 24bit, even 32bit or 96bit. Most humans can only recognize bit depth up to around 14 or 15bits, so its certainly arguable that anything over 16 bit is just overkill, causing larger files without any recognizable quality change. But, some claim to be able to hear this. Audio bit depth can be understood as the number of digits that can be used to represent your audio sample.

Confused? Does this graphic representation help? To make life simpler, lets take a 2-bit sample of this sine wave. 2 bits means we can represent 0, 1, 2, 3 (or 00, 01, 10, 11 in binary).

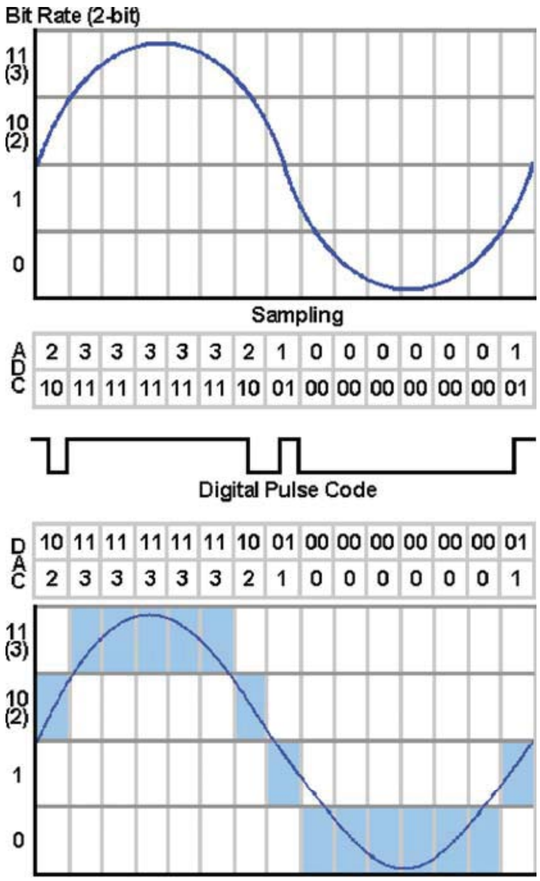

See how choppy the blue blocks look? This will result in an audio wave that should sound like a perfect sine wave, but rather will sound distorted.

In 8 bit resolution, 8 bits are used to store the measurement, so the measurement can only take 2^8 or 256 different values. In 16 bit resolution, 16 bits store the measurement, so the measurement can take on 2^{16} or 64,000 different values.

Sampling at 8 bit saves you some space, but often introduces a lot of "hiss" into a recording. Very roughly, 8-bit resolution is telephone quality, 16-bit is CD quality sampling and results in clear sound. This is a case where you'll need to consider your audience. What sounds terrible on your nice home stereo system, might sound just fine coming out of terrible speakers on your computer. Or, maybe not. You (or your client) will need to balance audio quality and file size considerations.

Seeing Sound

You've already seen some sound. Remember the most basic sound there is—the sine wave.

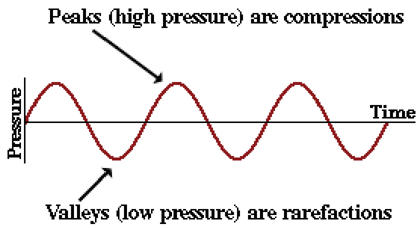

In Audacity, ProTools, or anytime you're working with digital audio, you'll typically be working with more complex waveforms than sine waves (although, sine waves are pretty cool). These complex waveforms that you'll run across while editing digital audio will look at bit more like this:

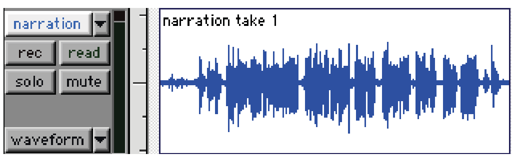

As you get experience editing audio, you'll learn to sort of "read" the waves, so just by looking you'll know where to perform edits. The horizontal scale (across) represents **time**, and the vertical scale (up and down) represents **amplitude**. So where there is very little amplitude, is a good indication of where you can seamlessly make an audio edit.

In some studios or other basic audio production class, you'll start by editing audio the good old fashion way—with reel-to-reel tape machines. You'll "scrub" the tape over the playhead, slowly moving the tape back and forth with your hands, until you find the silences between words or musical notes. That's where you'll want to make your edits - in the silence.

What you're typically looking for are spots where no one will hear that you've made an edit. If you're editing spoken word, you'll also need to consider intonation (or how the words are spoken), grammar (does what they say make sense anymore), and pops and clicks that are the by-product of poor audio edits. Good audio editors will also make sure the pacing stays consistent, and that narrators breaths are naturally spaced.

Pops and Clicks? How do you avoid getting pops and clicks? Edit at the decay or the end of the sound, and before the attack of the next word (or note if you are editing music).

ADSR
All sound has a certain **envelope** to it - or the way its amplitude changes over time. This is referred to as an

audio envelope or ADSR. ADSR stands for Attack, Decay, Sustain, Release of a sound. It's not the sound itself (which would be its timbre, frequency, etc), but rather a way of describing the intensity of the sound over time.

Let's take a piano note as an example. It may have an envelope that looks something like this:

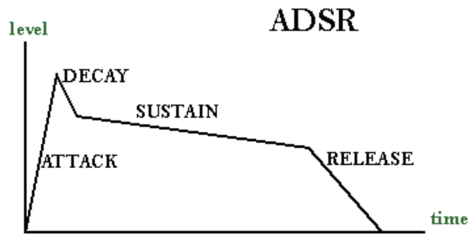

Most sounds have an initial, fairly quick **attack,** or beginning of the note. Think of the piano note. There's silence, then suddenly a loud sound. The loudness decreases (**decay**) a slight bit, then the sound seems to hold (**sustain**) as long as you hold the key down. **Release** the key, and the sound naturally fades away.
In the analog audio editing world, you'll need to be able to hear and really listen for the attack of a sound to make decent edits. In the digital world, you need to listen, but you also get to look to see where the sound starts.

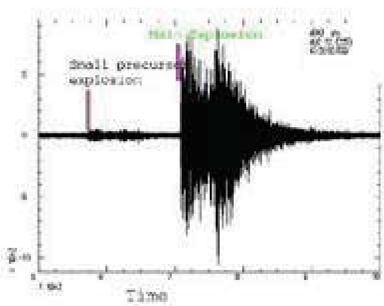

Soundcards

You can add a PCI sound card that goes right into your computer's motherboard. PCI (peripheral component interconnect) slots are standard buses on pretty much every computer. Your computer can have several of these slots which you can fill up with soundcards, graphics cards, or whatever add-on specialties you need.

In general, Macs usually come with all the cards you need for general purpose graphics, audio, and video built-in to the system. And, as we progress forward with technology, these internal PCI cards are quickly being replaced with external devices that typically connect to your computer with a USB or Firewire connector. Many Windows PCs let you choose your on-board soundcard.

Most sound cards will let you record audio into your computer, get a line-level output, a speaker or headphone output, and maybe even a digital audio or MIDI in/out port.

These sound cards interpret all the digital data (samples) that are created and they know how to convert the samples back to analog audio waves that you can hear. They are DAC and ADC, or digital to analog converters and analog to digital converters. They are really the best way to get audio in to your computer.

If you have a "headset" microphone that has mini-stereo plugs, you plug these into the soundcard jacks on your computer.

USB or Firewire Audio Interfaces

USB 2.0 and FireWire are both speedy, and efficient at handling lots of data. (USB 1.1 is not really fast enough to handle audio.) Nowadays you can get a whole range of audio "boxes" that you can use to record multiple tracks of very high quality audio directly to your computer's hard drive.

Very popular options are the Digidesign audio interfaces. Many schools have a ProTools|HD setup, which is quite nice, but out of the price range for the average home audio person. For us, there's the Digidesign MBox line, available in either in USB or FireWire depending on your needs. These devices are great for audio recording, and they can get you "in the game" for about $300 on the low end. I used the entry level MBox at my job at Apple recording professional voice-over talent, so don't always assume you need the "top of the line" device.

Recording Hardware

There's some interesting portable audio recording hardware out there that's very convenient. You can check out the Zoom PS04 that is a handheld 4-track recorder which records to a Compact Flash Card. A former student turned me on to these, and they seem pretty convenient. The benefits to a Compact Flash recording device are at least two-fold. One, you don't have to worry about a tape device that has working parts in it (ie physically moving tape over heads), and so they should be pretty durable. And, better still, you can transfer the "files" right from the device to your computer, just like copying a file (and NOT like having to play the tape back and waiting "real-time" for it to record to your hard drive).

Personally, I just needed two channels of recording, so I did some searching and found two others - one by mAudio that I liked (and CCSF Broadcasting has several of these). I personally opted for a portable Marantz, which is a bit larger in physical size, but has the added bonus of real XLR inputs, so you can easily use nice microphones with it if you want. So, for $100 more I opted for the Marantz PMD 660 and I've been using it to record interviews, either with its internal mics, or with external mics depending on how much set-up I want to do.

Anyhow, these Flash Media recorders are pretty nice - and perfect for podcasters (more on this later) and for recording without lugging a laptop around.

DSP cards
DSP cards, also known as "farm" cards, actually contain DSP (digital signal processor) chips to process specially-coded applications. Traditional audio cards facilitate audio I/O with your computer, but DSP cards go a step further by taking the processing load off your main computer's CPU. The result is greatly enhanced performance, and the ability to run more simultaneous tracks! They are becoming less necessary as CPU's become faster and more efficient, and more and more of the input is being handled by external

Firewire and USB dedicated boxes. Multi-core processor remove some pressure from needing to buy special DSP cards. Historically, these DSP cards are expensive, and do only one thing - process audio. So they are not very practical for most "general" media makers.

Speakers

Are you listening to (or worse, mastering audio) on your computer using your computer's built-in speaker? And, is that speaker located right in the computer case, situated someplace on the floor? Or coming out of tiny speakers on your laptop? This is no way to hear audio.
There are really great external (typically USB) reference monitors that you can attach to your computer to get awesome audio from your computer. It's really worth it to get some of these speakers, even if all you ever do is listen to CDs in the background while you work. They will let you hear what you record with more accurate detail.

ProTools

ProTools is the audio editing application that you'll find used in most of the other audio production courses in the Broadcasting Department. ProTools is an industry-leading audio editing application. You can have a "low end" system (ProTools LE running with an MBox) or you can go to a top-of-the-line ProTools|HD system.

Audacity

Audacity is a very good little audio editing application - and it has the benefits of being:

- free
- available for Mac and Windows systems
- fully functioning, never expires, open-source software

Audacity has a pretty similar interface to ProTools, and it's actually easier to use. If you learn Audacity now (popular with podcasters) you'll quickly be able to transfer your skills to ProTools when you take a more advanced course in audio production. But your basic editing skills learned here will certainly help when you come to use their nice equipment!

QuickTime Pro

For $30 you get a simple, but really powerful, audio and video **capture**, **editing**, and **authoring** tool - available for both Mac and Windows. Yep, you can use QuickTime Pro to both record and edit video and audio files. Powerful, but simple, you can cut, copy and paste to edit media. Meaning, you get a "cuts only" editor - so you can chop up audio and reorder stuff, but you can't do cross-fades, dissolves, or work very precisely with multiple "tracks".

QuickTime Pro is often called the Swiss-Army Knife of the multimedia world - maybe it's not the "best" tool for the job, but it is always nice to have around when you don't have something more powerful. It's completely cross-platform, so it works seamlessly on both Mac and PCs. And, it also lets you import and export nearly any file format - so you can convert files, compress audio and video, create movies from still images, etc. We'll get into this more later.

GarageBand (Mac)

GarageBand (Mac-only) is one of the easiest to use digital audio applications, and it comes with hundreds of sound effects, music loops, and background sounds you can mix with your own recording. It comes included with iLife.

Adobe Audition (Windows)
Adobe Audition (about $350) is a professional-level audio editor that lets your record, mix, edit and master audio. Formerly known as Cool Edit Pro.

Video editors as Audio editors
Bundled for free with your operating system, you have iMovie (Mac) or Movie Maker (Windows) available for you to use. Both are mostly considered for their video editing abilities, but hey, you can use them to edit audio-only too.

Remember, most media makers don't need to know every intricate detail of every application, they need only know how to get going with a project, and where to look for help! I only mention these apps here so you understand that you can use whatever program you want (or are forced to use) as an audio editor. They all adhere to the same basic principles.

MIDI and Sequencing

MIDI is simply a protocol (like a digital language). The MIDI spec was created back in the early 1980s, so electronic musical instruments could communicate with each other. MIDI stands for Musical Instrument Digital Interface, and is used for many different things, including musical sequencing, composition, and is even used in live shows to run lights. It's a standard created by the music industry, so all digital keyboards "speak" the same language. Your computer doesn't care if you're using a Korg or a Yamaha MIDI device....all it knows is it's receiving digital data in the form of a MIDI signal.

However, don’t confuse MIDI data with digital audio information. These MIDI cables do not send actual audio waves. Rather than, say, recording a trumpet by using a microphone connected to your computer (digital audio), with MIDI, you're saving the recipe, or set of directions about how the trumpet was played. So, MIDI files are tiny in comparison to digital audio files because they are sending instructions, not audio. MIDI sends instructions that let your computer know things like which instrument you were playing, which notes you played, in what order, and for how long, and how loudly each note was played.

Unlike digital audio, your computer is being sent instructions on how to play the MIDI file, and your computer needs to have some kind of MIDI device (or capability) to actually play the audio. QuickTime provides a software MIDI device, so you'll be able to hear MIDI files (you can download them from the Internet). iTunes and Windows Media Player also gives you a MIDI player.

MIDI data is transmitted using a binary number system. A serial interface (typically USB nowadays) translates musical actions into numbers and sends one at a time through a MIDI cable. The message might have "Note ON, MIDI channel 1, play the 55th note, at a velocity of 123."

MIDI is a whole different domain of audio and music.

Audio Mastering/CD Burning