Max and MSP

MSP is an addition to Max that provides signal generation and processing objects. It works entirely in the Macintosh, which gives you advantages and disadvantages.

Advantages:
• You can do whatever you want, not whatever marketing thought would sell.
• It doesn't cost a lot of money.
• It doesn't cost any money to change your mind about what kind of sound you want.
• When a new computer comes out, your patches won't have to be thrown away, they'll run better!

Disadvantages:
• Your typical MIDI instrument can run circles around a Mac (even a G5) when it comes to audio processing. So dedicated instruments will always have more channels, thicker chords, etc.
• There can be a tiny delay when sounds pass in and out of the computer.
• You have to make things work yourself. An empty patcher won't make a sound.

Warnings:
MSP is not entirely bug free. You are going to have some crashes and freeze-ups, especially when your patches get large. You must do three things to cope with this:
• Save often, especially when you are about to turn on audio
• If it starts to behave strangely when audio is not on, save and reboot.

You are going to have to study the tutorial (which is a PDF file on the computer- open the manuals alias to find it). But here are some basic concepts to get you started.

The DSP Status Window

MSP operations are Set up in the DSP status window, which is found under the options menu. There are several important settings in this window.

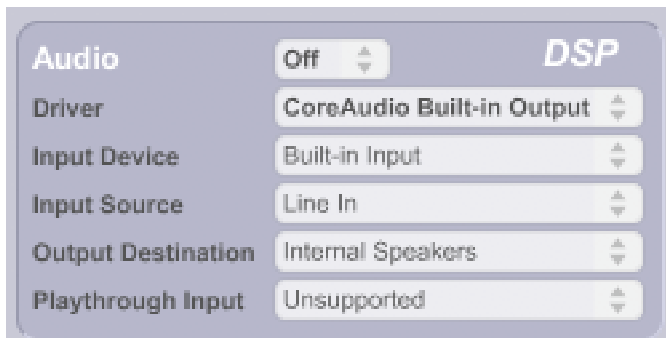

Figure 1.

The section shown in figure 1 allows you to select the sound interface and input source.

An interesting choice for driver is "nonRealTime". Which lets you compute complex processes and listen later, a la Csound. Some options will not fill in for some interfaces--these will have to set on the interface control panel or sound preferences.

Figure 2.

You can change the sample rate in the pane shown in figure 2.. I/O vector size is the number of bytes pulled off the interface at one time. This directly affect audio latency. The signal vector size is the number of bytes Max processes in a batch. These should both be as small as you can get away with. You know you aren't getting away with a setting when the audio pops and crackles or the user interface objects get sluggish.

The Max scheduler is the master clock that processes all non audio activity. When it is placed in overdrive, calculations triggered by midi or metro happen at interrupt level, without stopping for slow processes like screen redraws. In the old days, when computers were pokey, we kept this on but it doesn't make a lot of difference on Gigahertz machines. The Max scheduler is only accurate to within 2 milliseconds. If you want timing tighter than that, putting the scheduler in audio interrupt will get it down to once per vector. (Divide vector size into sample rate to figure out how often that is.)

Figure 3.

The pane of figure 3 tells how the computer is doing. Audio processing is pretty hungry, and if CPU usage gets above 70% or so, other functions of the computer, such as processing keyboard input will suffer.

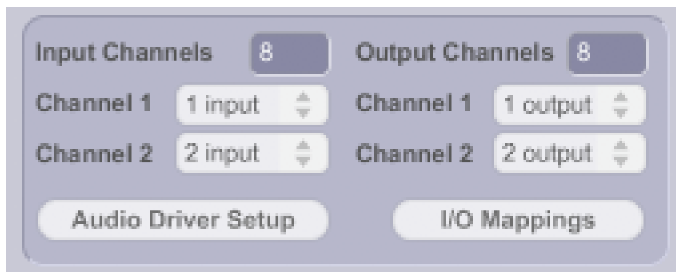

Figure 4.

Figure 4 shows where tochange inputs and outputs. Max's idea of channel 1 can be any input from the sound card. The I/O mappings brings up a window that lets you set all your interface supports.

MSP Basics

MSP objects have a tilde (~) after their name. Many of them have the same name (except for the tilde) as regular Max objects. Usually, the function is similar to their namesake.

Max works with messages- int, float, list, bang, that are sent once from one object to another. MSP works with signals that are flowing continuously whenever audio is on.

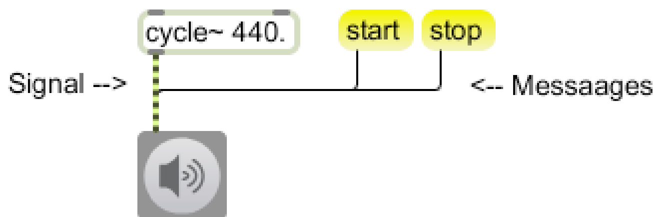

Figure 5.
Cords that pass signals are yellow with black stripes. The signals really consist of samples in batches called vectors. The number of samples in a vector can be changed for various reasons, but usually there are 64.

You can't watch a signal with a message box the way you can with standard Max messages. The best you can do is grab a single sample with sah~ (sample and hold), grab a bunch of numbers with capture~, or watch a display similar to an oscilloscope.

When a signal is captured, it looks like a lot of floats. The values should be between -1 and 1. A signal that swings all the way from –1 to 1 represents full scale and is very loud. Anything larger than this will distort if it gets to a dac~.

The most primitive signal comes from sig~, which provides a constant signal – 44100 copies of the same value per second[1].

Figure 6.
Number~ is another source of constant signal, with a user interface. It will also display the value of a signal that is connected to it, but only four times a second. That's only useful for slowly changing values. (The update rate and mode of number~ are set in the inspector.)

[1] Don't confuse a constant signal with a signal of constant amplitude. The latter is an audio signal that does not change in loudness. A constant is not audible, as its frequency (whatever the constant value) is 0.

Audio Output

The output of an MSP patch is the dac~ or ezdac~ (shown in figure 6, it looks like a button with a picture of a speaker on it.) To start audio, send either a 1 or the message "start". (You can just click on the ezdac~) This starts all audio in all windows. To hear only audio from one window, send "startwindow". Audio will stop with the stop message, or a. You can add signal objects while audio is running..

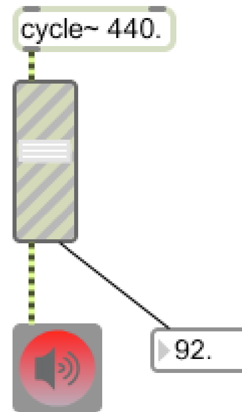

Figure 7.

The left and right inlets of the dac~ correspond to left and right stereo outputs. dac~ can have arguments that determine the output channels, up to 16.

Levels

To adjust the volume of a signal, you use either a multiplier [*~] (with a fractional value) or a gain~ slider. This looks very much like the slider but it faintly colored stripes.

The signal is applied to the left inlet and is available at the left outlet. An ints at the left inlet will move the slider. A change of 10 will adjust the level by 6 dB. The value 128 produces unity gain (no change), and the top position is 157. This setting can easily produce distortion. The position of the fader is reported at the right outlet.

The level does not change instantly when you move the slider. It ramps up or down so the signal doesn't pop. A float in the right inlet sets the ramp time in milliseconds. It defaults to 10 ms, which is usually fine.

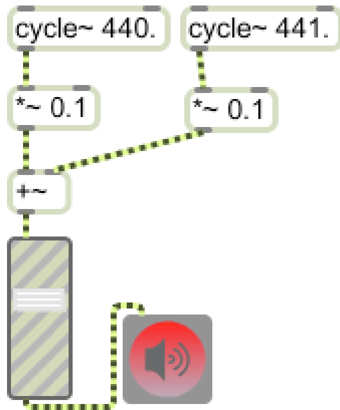

Figure 8
A multiplier is often a better way to control gain, if you don't need a user control. (You should still have one gain~ slider per sound.) If you think of the multiplier as equivalent to a VCA, you'll get the idea. I usually stick a multiplier after every signal producing element in order to get levels under control, as in figure 8. Notice that the adder [+~] object is used to combine signals. Think of it as a mixer. There's no difference between the two inlets -- they are both active continuously. Adder objects are not really necessary. Any signal inlet will add signals, but I find them a useful tool for organizing patches.

Believe me, even attenuating those cycles~ to 0.1 gives a plenty strong signal. Remember your decibels?

Db = 20 log v/V, where V is the peak output of cycle~.

20 log 0.1/1 = -20 dB
20 log 0.01/1 = -40 dB
20 log 0.001/1 = -60 dB

20 dB down is noticeably quieter, but when we add two signals like this, the peak values of the sum will hit 0.2, which is only one fifth of distortion level.

Audio input

Inputs are shown by the adc~ or ezadc~ objets.

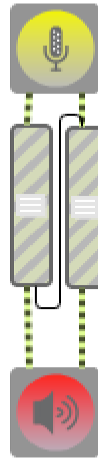

Figure 9.

Figure 9 simply passes two channels of audio through with gain control. Note that I have connected the right outlet of the left fader to the right fader. This is how you build a stereo control. The left fader controls both channels, and is all you would need to show in the presentation.

There are over 4 dozen MSP objects to process audio, many of which are described in other tutorials. Figure 10 shows a basic delay with feedback.

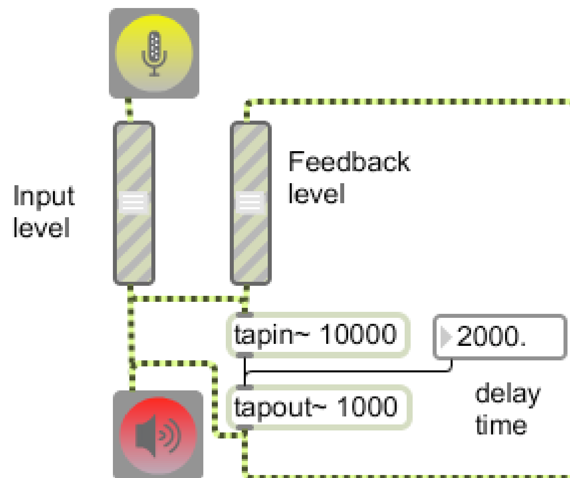

Figure 10.

Oscillators

The cycle~ object is the basic oscillator. When audio is on it puts out a sine wave at he indicated frequency.

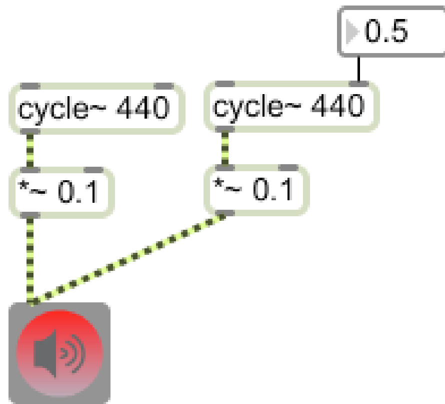

Figure 11.

There are two inlets to cycle~ the left one controls the frequency, the right one controls phase. Frequency is in hertz (and there's a neat object called mtof that converts midi note numbers to frequency.)

Phase is a fraction of the wavetable, so 0.25 is 90 degrees, 1.0 is a full 360. The patch at the left doesn't make any sound, because 0.5 is 180 degrees out of phase, and you know what that does!

For more about oscillators, see Basic Synthesis in MSP.

Playing Sound Files

There are two basic ways to play sound files. With buffer~ (and its associated objects) which plays from memory, and with sfplay~ which plays files from the hard drive. There are advantages and disadvantages to each. Playback from memory is going to be instantaneous, but the total playing time you get is limited by the memory available in your computer. Remember that a minute of sound takes 5 megabytes of memory, twice that for stereo. There is no limit to the length of a file played by sfplay~ but the number of files you can have playing at one time is limited by the speed of your hard drive system. Most systems will easily do 8 tracks.

Buffer~ and Friends

The buffer~ object is a holder for sound files. You have to give it a name and a size, then put some sound into it. The size is in milliseconds, and will be a two channel unless you specify how many in a third argument. Loading the sound is done with the read message. There are several variations:

Read	By itself read brings up a file dialog and you find the file in the usual manner.
Read filename	This will load in the file named… if the file is in the search path defined in the Max file Preferences option. Otherwise you have to specify the complete path, such as /Users/pqe/Documents/SFDIR/filename.aif

	Note that folders are separated with slashes. A slash at the beginning refers to the main drive.
Read filename offset	Offset is a number of milliseconds from the start of the file to begin reading.
Read filename offset duration	In addition, loads only part of the file.
Read filename offset duration channels	You can read only one channel of a two channel file if you want to.
Replace	reads a file and changes the buffer size to fit the file size.
Readagain	reads the last file, but with new options if desired.
Import	reads MP3 files.
Write	Write lets you save a buffer as a sound file.
Writeaiff, writesd2, writewave	These specify file type, saving a step.

If you double click on a buffer~ a little window pops up showing the sound file.

Play~

The play~ object plays whatever is in the buffer it points to (you have to specify a buffer name.) It's very primitive, rather like turning a record by hand-- you put a signal into it, and as the value of the signal changes, play~ produces sound. Usually the signal is derived from a line~ object, but you can use a phasor~ to play loops. All play actually does is convert the signal coming in, (which represents time in milliseconds) into a pointer to a sample to grab from buffer and sends that sample out. The help file gives a fine example.

Groove~

Groove~ is the most useful buffer~ player. It can be pretty confusing until you realize exactly what it does. It is also maintaining a pointer into the buffer~, but it has built in line functions to move the pointer for you. If you send a float to groove~ it will cue up the pointer, but nothing will happen until it gets a signal -- this signal determines the rate of play, so usually a sig~ 1.0 will do it. Changing the value from sig~ will change the rate of play-- it will even play backwards. After it plays once, you have to cue it to the beginning to get it to play again.

Groove~ has a loop mode, where it will recue itself. You can set the loop points.

Record~

Record~ will put audio into the buffer~. You need to specify which buffer to use (the set message lets you change this) and number of channels if more than 1. Audio signals connected to the inlet(s) will be recorded when a 1 is received in the left. Recording is stopped with a 0. There are two mode flags for record:

Append	when off, recording always starts at the beginning of the buffer. When on, recording starts where it last left off.
Loop	when loop is off, recording stops at the end of the buffer. When on, as the end of the buffer is reached, recording moves to the beginning of the buffer, overwriting what is already there.

There are inlets to set the start point and end points in the buffer. These require floats that specify time in milliseconds. Record~ does not get a direct connection to its buffer~ It has a signal output that sends a location pointer. This could be connected to a play~ to synchronize playback from another buffer. This lets you patch multitrack features or set up a time remaining indicator.

Figure 12 is an example of a buffer in action:

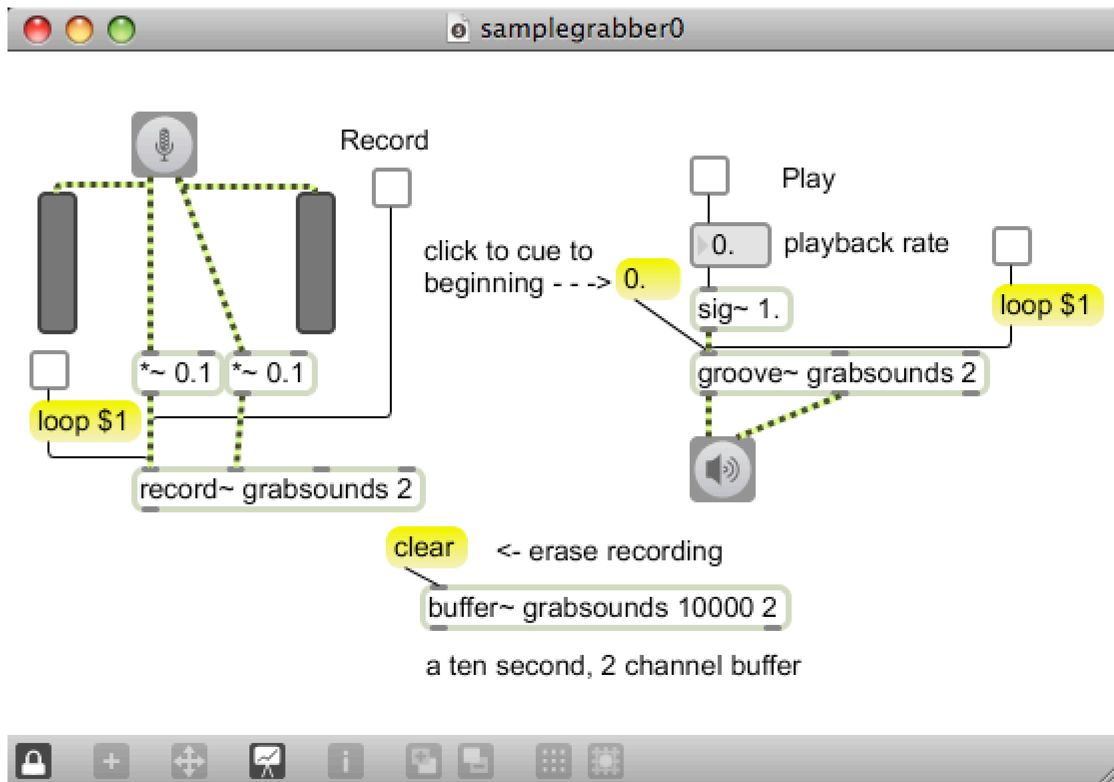

Figure 12.

Sfplay~

Sfplay~ plays sound files from the hard drive. The file can be in practically any format, although paired files (such as made by pro tools) will require two sfplay~s. Arguments to sfplay~ are:

Name of an sflist~ object ,	This is optional. If there is one, you can load several files and have them ready for instant playback.
Number of channels,	up to 8, apparently.
Size of play buffer	in milliseconds. A play buffer is necessary for smooth disk operations. If you put 0, the default size is used, which is usually fine. Adjust buffer size if you have slow disks that stutter when you play.
Number of position outlets	the first position outlet gives the time in milliseconds. This is rounded off from the actual sample locations. A second position outlet gives the round off error, so you add them to get the precise time. The help file shows how to convert this information into a time display.
Name	you can give the sfplay~ itself a name. Sfinfo~ can use this name to return information about the current file.

The reference on sfplay~ is 6 pages long, and it has extensive help files. The easy way to use sfplay~ is like figure 13.

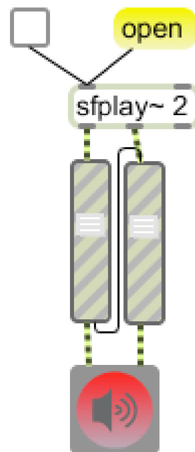

Figure 13.

This will play a file for you. If you want to use the file like a bank of samples, you can define locations as cue points with the preload command. Once a location has been defined as cue 3 for instance, a 3 (as opposed to a 1 which starts at the beginning) will start playback there. You can create and save complicated lists of cues with the sflist~ object. Sflist~ lets you have cues from more than one file. Since each cue has a buffer to start playback instantly while the disk catches up, there is a memory cost for using cues. Figure about 40k per channel.

Other commands like speed, seek, pause and resume make playback interesting.

Sfrecord~

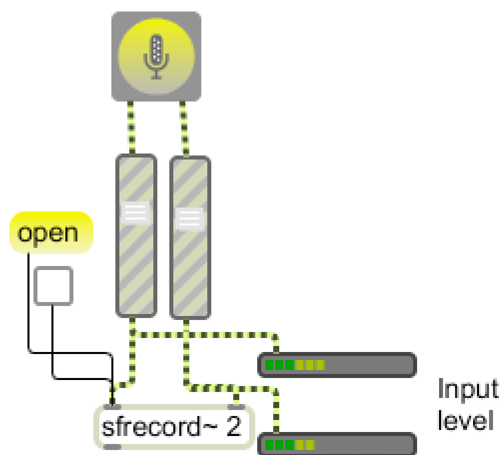

Figure 14.

Recording can be just as simple as playback. With a patch like the one above:

- Open a file to record into
- Send in a 1 or the message [record length] to begin recording.
- Send a 0 to stop.

There are options for various file formats and so on. I usually just use the sfrecord~ help file for my recordings.

Sfinfo~

Sfinfo~ provides information about selected audio files. Usually you want to know how long a recording is, and confirm that its sample rate is appropriate for current settings.

Testing Things

Sometimes the hardest part of working with msp patches is to find out what is going on. The tutorial shows several common ways of measuring signals- these are nicely illustrated in the help files.

Meter~ shows level. It flashes red if the signal goes above 1.0.

Avg~ gives the level whenever it is banged. You usually see a metro hooked up to it. There is also **average~** which does not require bangs and has several measurement modes.

Number~ periodically gives the value of the most recent sample at the right outlet. Like a sample and hold, it gives a funny pattern of numbers on audio, is most useful for slowly chaining signals, like the output of line~. When the arrow is showing, it creates a constant signal at the left outlet, which you set with the mouse.

Snapshot~ gives much the same thing, but reporting can be controlled by your patch

Capture~ lets you see a chunk of signal as a series of values. Tedious to use, but enlightening when everything else fails.

Scope~ Gives you pictures of signals. Its not much like an oscilloscope though, because what it really does is flash pictures on the screen at a rather slow rate. The help files and the tutorial don't really make the operation of scope~ very clear, so I'll have a stab at it.

The scope~ object captures the signal into a series of buffers that will be displayed on the screen. You can set the number of buffers by sending an int into the right inlet. You can set the number of samples per buffer at the left inlet. When the buffers are full, each is shown as a line segment from the value of the first to last sample in the buffer. (With the default display size, each buffer gets a single pixel in width, so the lines go straight up. If you stretch the display, you can see some slant.)

The total number of samples that will be on the screen is the product of number of segments per screen and samples per buffer. The defaults are 128 x 128 (16384) or about 1/3rd of a second at 44.1 kHz. This is suitable for showing low frequency waveforms, but the screen sort of jumps, and if you increase the frequency to the audio range, the display will probably just give you a line. To get an image of a high frequency tone, we need to display fewer samples. It's usually prettiest to keep 128 buffers and go to fewer samples per buffer. I usually start with around 8. If there are too many cycles showing, reduce the number. If you only see part of the wave, increase the number.

For high frequencies, reduce the number of segments per screen until you get a stable waveform. It may look strange, but then again the waveforms digital systems produce at high frequencies really do look like that.

With real oscilloscopes, you can set the sweep rate to a particular time per screen marking and use triggering circuitry to make the waveform sit still. The settings for scope~ are restricted to integer sample numbers, so we usually can't get convenient time settings. We can however, calculate the time per division we are seeing. This patcher shows how.

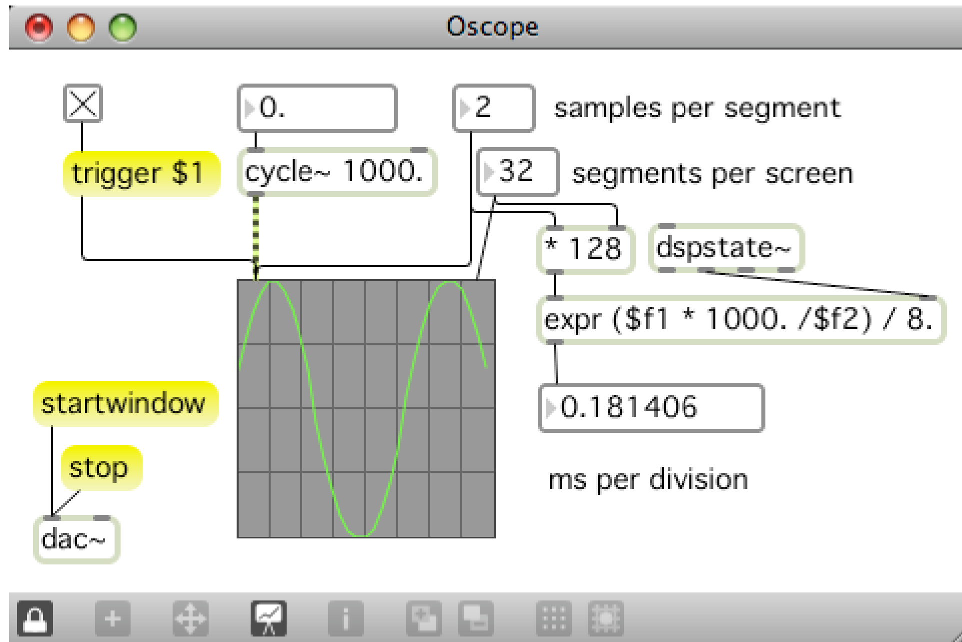

Dspstate~ gives the sample rate. With 8 divisions per screen and (samples per segment * segments per screen) samples on the screen, the expression object shown will give ms per division.

The trigger 1 message to the scope will stabilize the display if the other settings are in the right range.

How it all fits together

Although listening to audio and watching the screen gives the impression that everything is steadily chugging along, the system is really backing up and stretching out like cars on a crowded interstate.

There can only be one sample rate in effect at a time. This is determined by the sound card settings. Assume it's 44.1k for this exercise. That means the card needs a sample (X the number of channels) every 0.0226 ms. To make sure there's always something there, the program can run ahead of the card, up to whatever the output buffer size is, 1024 on the one I use.

The MSP audio interrupt is loosely locked to the audio card. The audio interrupt runs at (Sample rate)/ (signal vector size) . If the vector is 64, an audio interrupt happens approximately every 1.45 ms. The duration of the audio interrupt depends on the amount of processing necessary for every audio object to do its thing on one vector. Naturally, if it takes longer than 1.45 ms, the output buffer will start to drain. You can get away with a couple of long vector periods (like when something turns on or goes the long way around a branch), but they better be balanced by short ones. If the audio processes get ahead of the card, the card will say so, and an interrupt is skipped.

Note that there is a buffer for everything that is contributing signals to the system. The card input and any files being read all put their data into a buffer, and on the audio interrupt, a vector's worth of samples are taken out. If the buffers are too big, input to output takes too long, but if they too small, audio processing is broken up. (The buffer~ object is an extreme example of this, as it contains the entire audio recording.)

Also note that for proper playback, any recording should be played at the sample rate in use when it was made. If the current sample rate is different, the program has to do some interpolation to compensate. MSP is doing this anyway for variable rate playback. When you specify sample rates for recordings, you are telling MSP the original SR so it knows what to do. Sample rates are marked in audio files, but it's easy to get them messed up in some programs.

When processing time gets tight, it would be nice to skip some calculations, especially if they are just giving the same number over and again. You have to do this in a closed off part of the system that is running at half the sample rate. Any signals entering this area have to be decimated (downsampled) on the way in, and resampled back up on the way out. The poly~ object is a nice enclosed area where this is possible, and often necessary. In a leisurely patch, some things may sound prettier if run at twice the sample rate. I'll be interested in finding out what these are.

That's the bottom layer- Next we find the old Max scheduler. It can be run by the Mac internal clock (which is a bit pokey and the source of years of complaints) or, for accuracy, can be tied to the audio interrupts. The scheduler ticks over once a millisecond. Note that this is faster than the audio interrupt, but on most ticks it has nothing to do. When there is something, such as a metro to bang, everything attached starts happening in a long involved chain. If it's not finished by the next audio interrupt, the CPU puts a bookmark in and does the audio chores. That's why it's called an interrupt. When the audio chain is through, scheduled tasks resume. (The pecking order of the scheduler and audio can be reversed with the prioritize MIDI option)

There's an even slower layer. Every once in a while, after scheduled tasks are done, the program asks the event manager (part of the operating system) if the user has typed a key or clicked or something. If so, another long processing chain is kicked off, which can be interrupted by audio, or if overdrive is on, scheduler tasks and MIDI input. After this sort of thing is done, Max relaxes for a moment to let the OS catch up on outside work like sending you messages about your network. This is all interruptible, which is why the system will lock up if you pile on too much audio.

Max also keeps a list of things to do when the cows are milked and the chickens fed. These are things like file operations and graphics drawing. You have often seen the effects of this on number boxes and the like. You can move any part of a patch into this zone with the defer object.

This is beautiful system, but it makes some things difficult and unpredictable. Here are some questions to ponder:
What if a scheduler task needs to know the current signal value? Which of the 64 numbers in the current vector does this mean?
How is the operation of graphic sliders affected by audio load?
Why the sig~ object?
Why doesn't Signal Scope behave like the one on my bench?
How can midi controllers produce smooth audio effects?