MSPeak

A glossary of dsp terms.

First a quick review of what is going on. Digital audio is focused on sending numbers to a digital to analog converter in a manner appropriate to hearing the result as sound. There are several scenarios possible:

1. The digital system is merely a transmission medium. That is, numbers are sent out just as they came in with no deliberate transformations. Think telephone. For good performance, it is important that the rate and precision of the output match that of the input., and that no values are accidentally changed.

2. The digital system is used for storage. Numbers are moved to a more or less permanent storage device like a hard drive or a CD. Later, they are read off the device and sent out. In addition to the concerns for transmission, we worry about loosing any of the data because of defects in the media. We also have to worry about synchronization, reconciling the needs and ability of the storage devices with the necessity for a steady stream of data.

3. The audio is changed somehow. Generally we want changes that make musical sense, often mimicking techniques we have gotten used to in analog processes.

4. The audio is made up as we go along. Now we have to face the problems of managing an enormous number of parameters in a convincing way.

All of this has to be managed in the context of a computer system that can only do one thing at a time. These things are limited to moving numbers around and performing basic math operations. Some math operations are quick, others cumbersome, and unrelated demands on the computer's attention arrive at unpredictable times. To keep things moving smoothly, the numbers are moved in and out of pools of data called buffers - the buffers can fill with the arrival of a block of data from a hard drive, then slowly drain as samples are removed and processed. Other buffers receive the processed data, where a few samples await their moment for playout.

Vectors and interrupts

Since CPUs usually spend more time setting up to do something than actually doing it, the audio samples are processed in batches called vectors. In MSP, the vectors are passed along the yellow patch cords. There is a timer running that produces something called audio interrupts. When the interrupt happens, each MSP object processes one vector of audio. There's another timer called the Max scheduler. It is much slower, ticking about a thousand times a second. Regular messages are passed and processed when the scheduler ticks over. (When overdrive is turned on, arriving MIDI messages are processed immediately on an interrupt of their own.)

Making waves

To generate periodic waveforms, we load one cycle's worth of the wave into a buffer called the wavetable. Then we step through the wavetable, moving a pointer one location on each sample time. When the end of the table is reached, the pointer is wrapped around to the beginning of the table. This will be heard as playing the waveform at the original pitch.

To get different pitches, we move the pointer differently. We add a sample increment to the pointer each time. If the increment is less than one, the resultant pitch will be lower than the original. If he sampling increment is greater than one, the pitch will be higher. The formula for the frequency is:

F = (sample rate) * (sample increment)/ (size of table)

The math to make this happen is most efficient if the wavetable size is a power of 2. When low pitches are played, the waveform will be distorted, because we are really playing samples twice. To get a smoother waveform, we make up samples when the sample desired falls in the crack between two existing samples, a process called interpolation. Interpolation takes a bit longer, so it will limit the number of waves you can generate at once. Sometimes it's better to skip the interpolation and just use large wavetables. Interpolation algorithms often require that the wavetable be a power of two plus one. (That extra spot holds a copy of the first sample.)

Processing

Changing the amplitudes of signals is done by multiplying the sample values by a scaling factor. Sometimes the scaling factors are taken from another signal, which may change relatively slowly to provide an envelope.

Believe it or not, the only other process available in a digital system is delay. Delay is available by writing the signal into a buffer and reading it out a little later. Since you can read form more than one location in the buffer at once, various delays can be combined. Filter design is the artful combination of delays. If the delay time is the same as the period of the signal, the delayed samples will be in phase with the signal and amplify it. If the delay time is half the period of the signal, the delay is out of phase, and will reduce the output. How much amplification you get can be controlled by scaling the delayed values. By adding multiple delays, you can get just about any kind of filtering.

Delays can also produce reverb and the usual list of echo effects.

Here's a glossary of the terms we encounter often in digital signal processing:

ADC--

Analog to Digital Conversion. The essential first step, a transformation of sound pressure into a list of numbers. The accuracy of the transformation is heavily dependent on the number of measurements taken per second (sample rate) and the number of bits available to represent the waveform. (word size).

AIFF

Audio Interchange File Format. Apple's native audio format. When Macintoshes started doing audio, there was no standard format, each program used its own. After a while, most applications support AIFF, although it's often second choice. A few years ago Microsoft declared their own standard, WAV. The differences in file formats are usually insignificant -- the audio is the same, but the header lists parameters in a different order. The big exception is SD2, which has left and right sample data alternating, which means if you play an SD2 file with AIFF settings, the pitch will be an octave low.

Alias

The error you get when trying to reproduce a frequency greater than half the sampling frequency. A common problem in synthesis as you reach for high notes, and is inherent in trying to synthesize waveforms that have high partials. In Msp it's pretty audible when you play the phasor~.

Allpass

A filter that does not change the amplitude of the signal, but changes phase of selected frequencies. If a signal changes level, allpass will briefly color the sound. Allpass mixed with original produces flanging.

Amplitude

The height of a graphed signal. Related to loudness. Digital systems adjust amplitude by multiplying sample values by a scaling factor.

Array

In computer systems, a section of memory that can be accessed as if it were arranged in rows and columns. It's handy for quickly finding associated sets of values.

Audio interrupt

A moment in time when non audio chores are suspended and the computer concentrates on signal processing. Audio interrupts must happen fast enough for the computer to keep ahead of the output sample rate.

Band-limiting

Excluding audio outside of desired frequency ranges. Bandpass filtering merely reduces out of band signals.

Biquad

This is a filter. It's actually an imitation of an analog filter, and is very basically implemented in msp. Figuring out the coefficients for biquad~ takes some pretty heavy math, so filtergraph is provided to do it for you. Usually, I use filtergraph to get the sound I want, then copy those coefficients into the biquad~ in the final patch.

Buffer

A section of memory that holds audio data. Buffering is necessary to allow for subroutines that process data in bursts or at different rates to keep up with each other. It's also necessary for hard drive access because drive heads have to move unpredictable distances across the disk to find the data. Buffer~ is a copy of an audio recording that is kept in memory for quick access.

Carrier

This is an old radio term, and refers to a steady signal that will modulated with some kind of information. It's used in FM synthesis.

Cartesian

Refers to the normal system of graphing on square grids. A waveform is an example of a Cartesian function. Cartesian coordinates have x, y and z components. The point where $x = y = z = 0$ is the origin. The line where $y = 0$ and $z = 0$ is the x axis and so on. In audio, the z direction is often replaced by an imaginary axis to represent coefficients of the square root of -1.

Clipping

A waveform gets cut off at the top if its amplitude is too large for the size of the data words that are representing audio in the system.

Clock

Any circuit that provides a steady time reference. The 44.1 kHz for audio comes from a clock.

Convolution

Technique in which the points in one waveform are multiplied by <u>all</u> of the points in another. Hopefully one of the waveforms is pretty short. Say h(n) has four samples. Each point in the convolution f(n) * h(n) is equal to
$f(x) * h(0) + f(x) * h(1) + f(x) * h(2) + f(x) * h(3)$.
If h(n) represents a filter impulse response curve, a convolution like this will give a filtered output. It's all a mathematical way of describing the effect of filtering by summing delays.

Cosine wave

A waveform that describes the simple rotation around an angle. The numbers could be taken straight from a table that gives the cosine of evenly spaced angles. A cosine wave is just like a sine wave, but it starts at 0, avoiding a pop.

DAC
Digital to Analog converter. It provides the actual sound of what we are computing. Luckily, building them is somebody else's problem.

dB
Decibels - you had better already know this. In the digital world, multiplying the sample values by 2 adds 6 dB to the amplitude. Another way of looking at is multiplying by 0.1 reduces the signal by 20 dB.

DC offset
A constant value that is added to all samples. Typically, the midpoint of waveforms is 0 volts. If there is a DC offset, the values swing further one way than another, not usually a good thing. DC offsets can come from hardware or software problems.

Defer
Connect a Max process at the low priority. See the discussion of overdrive. Messages that are deferred will be processed at the front of the low priority queue, anything sent to deferlow will be the very last thing to happen.

Delta
Any sort of change from one moment to another.

DFT
Discrete Fourier transform - the algorithm that performs Fourier transforms on sampled data.

FFt
Fast Fourier Transform. An algorithm that is sped up by making some assumptions about the waveform analyzed, namely that only harmonics of the frequency implied by the size of the waveform table are present.

Filter terms
A filter is a circuit or software routine that changes the frequency response of a signal. The terms you really need to know are:

- Frequency response - this is a graph of the output amplitude produced by applying a sweep of frequencies to the filter..
- Phase response- this is a graph of how the phase of signals is changed at various frequencies.
- Passband-- the range of frequency that is relatively unaffected by the filter.
- Cutoff frequency -- the frequency at which the output is reduced 3 dB below the level of the passband.
- Bandwidth- a band pass filter has two cutoff frequencies. The bandwidth is the distance between them.
- Slope -- the rate at which output declines outside of the passband.
- Q- for band pass filters, the ratio of amplitude to bandwidth. High Q implies a narrow peak.

Foldover

A type of digital error where if values exceed the maximum they jump to the minimum negative value instead of clipping.

Fourier Transform

The technique of changing the representation of a signal from the time domain to the frequency domain or back. The time domain representation is the familiar waveform. The frequency domain representation is a series of snapshots of the spectrum. Each snapshot (often called a frame) has two parts, one representing amplitude, the other representing phase. The numbers in the snapshot, which are called bins, represent the amplitude at each frequency, which is some multiple of (sampling frequency)/ (number of bins). If the transformation is done accurately, there is no loss of information going either way. The reason to transform is that some operations are easier to do on a frequency domain signal.

Frequency

The number of times a waveform repeats in a second. Used to be CPS, now measured in Hertz (Hz) and kilohertz (kHz).

Gain

The amplitude change in a signal through a circuit. It's a ration between the input amplitude and the output amplitude. A gain of 1.0 (unity) means no change. To convert gain to dB use the formula

$$20 \log G$$

To convert dB to gain:

$$10^{(dB/20)} \quad \text{or} \quad pow(10,dB/20)$$

IFFT

Inverse Fast Fourier Transform. An algorithm to convert a signal that is in FFT form back to a waveform.

Impulse

A digital signal that consists of a 1.0 followed by 0s. This is used in the design of filters and reverbs. Impulse response is a waveform that is produced by an impulse signal applied to a program or system. (Think of firing a gun in a church.)

Interpolation

Finding numbers between numbers you know. When you increase the sample rate of an existing wave, you use interpolation to calculate new values between the existing ones.

Inversion

Changing the sign (or otherwise complementing) all of the values in a waveform. In hardware, this can usually be done simply by exchanging the two wires that carry a signal.

Latency

Delay time, generally defined as the time it takes a system to convert the input from analog to digital, process it, and convert it back to analog.

Lookup

In general, using a number to find another number in some kind of table. We generally use lookup to calculate a sample value in a waveform synthesizer, apply distortion to a signal, or find the phone number of tech support.

Mixing

Combining two or more signals. This is accomplished by adding the samples in two signal streams. Some mixing systems include other features, such as adjusting the amplitude of the input signals.

Normalize

Adjusting the amplitude of a file so that the highest sample value is the maximum defined in the system. In a system with 1.0 representing full scale, this is done by multiplying all of the samples in the file by 1/highest.

Overdrive

In Max, overdrive is the direct connection of MIDI events to any associated objects, bypassing the normal "one thing at a time" schedule. When overdrive is on, an arriving MIDI event interrupts any other processing that may have been going on. When overdrive is off, MIDI events are queued up and wait their turn. This can make response to playing seem sluggish. On the other hand, when overdrive is on, mouse actions may seem slow when a lot of MIDI is coming in. In MSP, the highest priority is always given to the processing of audio, which can delay both MIDI and mouse actions. Metro and other timed actions typically happen at the lower priority (hence are interruptible by MIDI actions), but can be raised to the audio level of priority or the MIDI level of priority when "scheduler in interrupt" or "scheduler in overdrive" are set in the DSP options menu. Jitter message passing is always at low priority, but some jitter objects can hog the CPU and bring the system to a halt entirely.

Period

The time it takes a wave to cycle back to the beginning. Period, usually expressed in milliseconds, is the inverse of frequency.

Phase angle

Two different points along a cyclical waveform differ by phase. This can be considered an angle, because all cyclical action is related to turning things. (Think about phases of the moon, for example.) We generally speak of phase in degrees, but actually do the math in radians. In Max/MSP the phase of various oscillator objects is controlled by a signal from 0.0 to 1.0.

Phase shift

Any change in the phase of a signal, usually accomplished by some kind of delay. If you are shifting a cyclical wave or the phase shift occurs in a circuit that has an associated frequency (like a filter), the shift can be expressed as an angle, but it is usually specified in milliseconds or samples. A phase shift can be plotted as a graph against frequency to give a **phase response** of a system. Inverted signals are often described as being 180° out of phase, but this is only true for sinusoids and other symmetrical waveforms.

Phasor

A phasor is the subroutine that computes the phase angle for each sample in a wavetable synthesis program. In max/MSP the phasor~ object produces a ramp signal from 0.0 to 1.0.

Polar

A polar plot is any graph where a quantity is compared to an angle. We are most familiar with them from microphone directionality plots, but they have many uses in computer music, from the design of filters to decoding amplitude values from an fft.

Pole

In a hardware filter, a pole is a circuit than has a time based response, usually a resistor in conjunction with a coil or capacitor. In digital filter design, a pole is a delay unit, which models pretty much the same thing. It is the nature of the universe that a pole gives a 6 dB per octave filter effect, so a 24 dB per octave filter probably has 4 poles.

Queue

This is British for a line. (As in waiting to get on a bus.) In computer code, a queue is a list of values that will be handled in order as the CPU finds the time.

Sample

A single value in a list of values that represent a signal. Also used in synthesis to refer to a short recording that is played back when a key is pressed.

Sample Rate

The number of times a second a data acquisition device measures an input signal.

Scaling

Changing the value of samples by multiplication. Scaling by 0.5 reduces the signal to half its amplitude.

Sidebands
In various modulation processes that involve a carrier tone and a modulating signal, the resulting spectrum includes new components that are spaced above and below the carrier frequency at an interval equal to the modulation frequency. These are sidebands.

Sinusoid
A signal that is shaped like a sine wave. Both sine and cosine functions produce sinusoids. The Fourier process defines curves as mixes of sinusoids.

Signal
The changing voltage or stream of numbers that carries the information we are interested in. In audio, the signal is what is intended to be heard.

Vector
In MSP (and most other audio software) the audio stream is broken into short chunks of samples for efficient processing. The chunk is called a vector. If the vector size is 32, 32 samples are done in a batch. The down side is you have to fill a vector before you can do any processing, and you have to process a full vector before you can output anything, so you are always at least two vectors behind. With vectors of 1024, that's a 2 ms delay. (Plus whatever the process takes.)

Waveshaping
A technique of distorting a signal by using the sample values to look up an output value in a table. If the table contains a simple diagonal line, the output matches the input, but if the line is curved, various types of distortion will occur.

Wavetable
A section of memory that contains the samples for one cycle of a signal that is being synthesized. It's usually a power of 2 or a power of 2 plus one long.

Wraparound
In general, the effect of something that moves off one end appearing at the other. In graphics, a spaceship may fly off the left and reappear at the right, implying the screen was wrapped onto a cylinder. Another definition is the error you get when a binary number becomes too big for the register space (like 17 bits in a 16 bit system). Only the lower bits will survive.