Introduction to Computer Music

Week 10 Acoustics and Perception

Version 1, 24 Oct 2018

Roger B. Dannenberg

Topics Discussed: **Acoustics, Pitch Perception, Loudness Perception, Localization Perception, Audio Effects and Reverberation in Nyquist**

1 Introduction to Acoustics and Perception

Acoustics and perception are very different subjects: acoustics being about the physics of sound and perception being about our sense and cognitive processing of sound. Though these are very different, both are very important for computer sound and music generation. In sound generation, we generally strive to create sounds that are similar to those in the real world, sometimes using models of how sound is produced by acoustic instruments. Alternatively, we may try to create sound that will cause a certain aural impression. Sometimes, an understanding of physics and perception can inspire sounds that have no basis in reality.

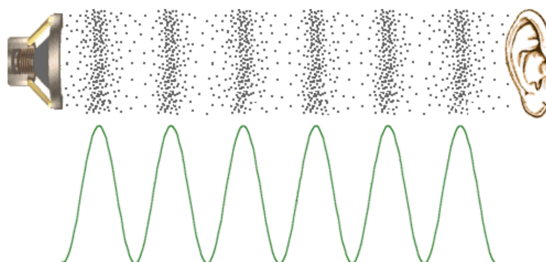

Figure 1: Sound is vibration or air pressure fluctuations. From mediacollege.com.

Sound is vibration or air pressure fluctuations (see Figure 1). We can hear small pressure variations, e.g. 0.001 psi (lbs/in^2) for loud sound. (We are purposefully using English units of pounds and inches because if you inflate a bicycle tire or check tires on an automobile – at least in the U.S. – you might have some idea of what that means.) One psi $\cong$ 6895 Pascal (Pa), so 0.001 psi is about 7 Pascal. At sea level, air pressure is 14.7 pounds per square inch, while the cabin pressure in an airplane is about 11.5 psi, so 0.001 (and remember that is a *loud* sound) is a tiny tiny fraction of the nominal constant pressure around us. Changes in air pressure deflect our ear drum. The amplitude of deflection of ear drum is about diameter of hydrogen atom for the softest sounds. So we indeed have a extremely sensitive ears!

What can we hear? The frequencies that we hear range over three orders of magnitude from about 20 to 20 kHz. As we grow older and we are exposed to loud sounds, our high frequency hearing is impaired, and so the actual range is typically less.

Our range of hearing, in terms of loudness, is about 120 dB, which is measured from threshold of hearing to threshold of pain (discomfort from loud sounds). In practical terms, our dynamic range is actually limited and often determined by background noise. Listen to your surroundings right now and listen to what you can hear. Anything you hear is likely to *mask* even softer sounds, limiting your ability to hear them and reducing the effective dynamic range of sounds you can hear.

We are very sensitive to the amplitude spectrum, Figure 2 shows a spectral view that covers our range of frequency, range of amplitude, and suggests that the shape of the spectrum is something to which we are sensitive.

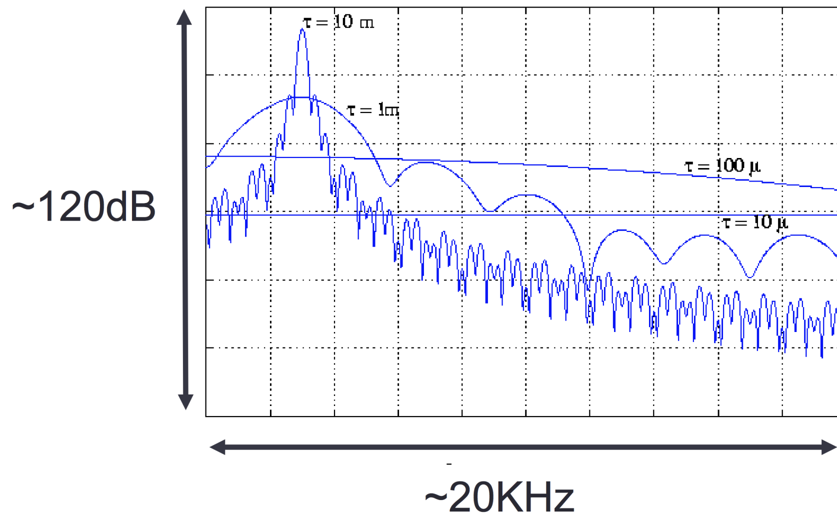

Figure 2: The actual spectrum shown here is not important, but the graph covers the range of audible frequencies (about 20KHz) and our range of amplitudes (about 120dB). We are very sensitive to the shape of the spectrum.

Real-world sounds are complex. We have seen many synthesis algorithms that produce very clean, simple, specific spectra and waveforms. These *can* be musically interesting, but they are not characteristic of sounds in the "real world." This is important; if you want to synthesize musical sounds that are pleasing to the ear, it is important to know that, for example, real-world sounds are not simple sinusoids, they do not have constant amplitude, and so on.

Let's consider some "real-world" sounds. First, noisy sounds, such as the sound of "shhhh," tend to be broadband, meaning they contain energy at almost all frequencies, and the amount of energy in any particular frequency, or the amplitude at any particular time, is random. The overall spectrum of noisy sounds looks like the top spectrum in Figure 3.

Percussion sounds on the other hand, such as a thump, bell clang, ping, or knock, tend to have resonances, so the energy is more concentrated around certain frequencies. The middle picture of Figure 3 gives the general spectral appearance of such sounds. Here, you see resonances at different frequencies. Each resonance produces an exponentially decaying sinusoid. The decay causes the spectral peak to be wider than that of a steady sinusoid. In general, the faster the decay, the wider the peak. It should also be noted that, depending on the material, there can be non-linear coupling between modes. In that case, the simple model of independent sinusoids with exponential decay is not exact or complete (but even then, this can be a good approximation.)

Figure 4 shows some modes of vibration in a guitar body. This is the top of an acoustic guitar, and each picture illustrates how the guitar top plate flexes in each mode. At the higher frequencies (e.g. "j" in Figure 4), you can see patches of the guitar plate move up while neighboring patches are move down. If you tap a guitar body, you will "excite" these different modes and get a collection of decaying sinusoids. (When you play a guitar, of course, you are strumming strings that have a different set of modes of vibration that give a clearer sense of pitch.)

Pitched sounds are often called *tones* and tend to have harmonically related sinusoids. For example, an "ideal" string has characteristic frequencies that form a harmonic series, as illustrated in Figure 5. This can be seen as a special case of vibrating objects in general, such as the guitar body (Figure 4.) The vibrating string just happens to have harmonically related mode frequencies.

We know from previous discussions that if the signal is purely periodic, then it has harmonically related sinusoids. In other words, any periodic signal can be decomposed into a sum of sinusoids that are all multiples of some fundamental frequency. *In physical systems, periodicity is characteristic of some kind of mechanical oscillator that*

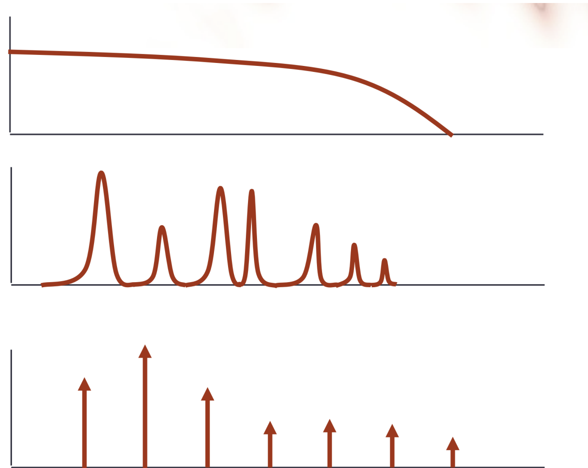

Figure 3: Some characteristic spectra. Can you identify them? Horizontal axis is frequency. Vertical axis is magnitude (amplitude). The top represents noise, with energy at all frequencies. The middle spectrum represents a percussion sound with resonant or “characteristic” frequencies. This might be a bell sound. The bottom spectrum represents a musical tone, which has many harmonics that are multiples of the first harmonic, or *fundamental*.

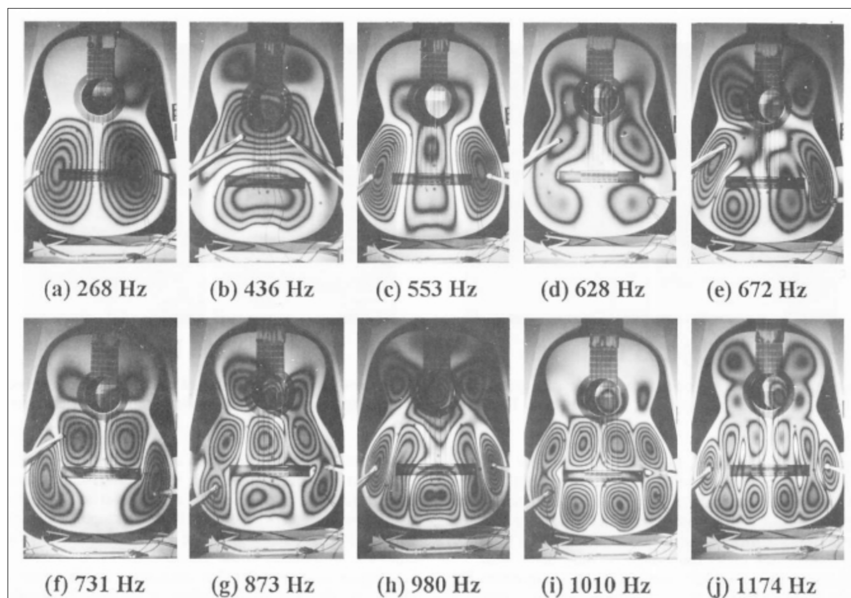

Figure 4: Modes of vibration in an acoustic guitar body. (From http://michaelmesser.proboards.com/thread/7581/resonator-cone-physics)

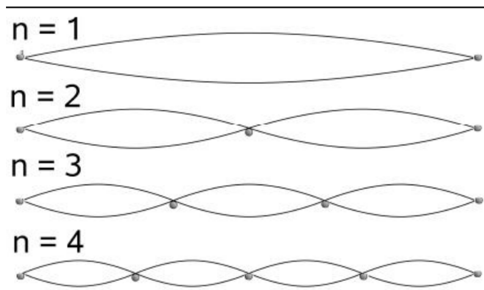

Figure 5: Modes of vibration in a stretched string. In an "ideal" string, the characteristic frequencies are multiples of the frequency of the first mode. The modes of vibration are independent, so the "sound" of the string is formed by the superposition of all the vibrating modes, resulting in a harmonic spectrum. (From phycomp.technion.ac.il)

is driven by an outside energy source. If you bow a string, sing a tone, or blow into a clarinet, you drive an oscillator with a constant source of energy, and almost invariably you end up with a stable periodic oscillation.

In summary, there are two ways to obtain harmonic or near-harmonic partials that are found in musical tones: One is to strike or pluck an object that has modes of vibration that just happen to resonate at harmonic frequencies. The other is to drive oscillation with a steady input of energy, as in a bowed string or wind instrument, giving rise to a periodic signal. In addition to these sounds called "tones," we have inharmonic spectra and noisy spectra.

2 Perception: Pitch, Loudness and Localization

2.1 Pitch Perception

Pitch is fundamental to most music. Where does it come from? How sensitive are ears to pitch? Our sense of pitch is strongly related to frequency. Higher frequencies mean higher pitch. Our sense of pitch is enhanced by harmonics partials. In fact, the connection is so strong that we are unable to hear individual partials. In most cases, we collect all of these partials into a single tone that we percieve as a single pitch, that of the lowest partial.

Pitch perception is approximately logarithmic, meaning that when pitch doubles, you hear the pitch interval of one octave. When it doubles again, you hear the same interval even though now the frequency is four times as high. If we divide the octave (factor of 2 in frequency) into 12 log-spaced intervals, we get what are called musical semitones or half-steps. This arrangement of 12 equal ratios per octave (each one $\sqrt[12]{2}$) is the basis for most Western music, as shown in the keyboard (Figure 6).

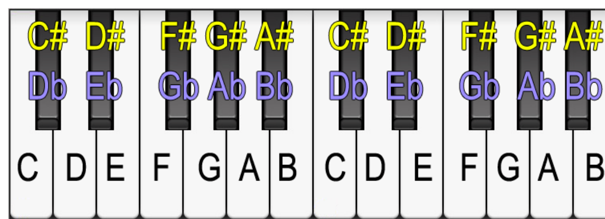

Figure 6: A piano keyboard. The ratio between frequencies of adjacent keys is $\sqrt[12]{2}$. (From http://alijamieson.co.uk/2017/12/03/describing-relationship-two-notes/)

We can divide a semitone ($\sqrt[12]{2}$) into 100 log-spaced frequency intervals called cents. Often cents are used in studies of tuning and intonation. We are sensitive to about 5 cents, which is a ratio of about 0.3%. No wonder it is so hard for musicians to play in tune!

2.2 Loudness

The term *pitch* refers to *perception* while *frequency* refers to *objective* or *physical* repetition rates. In a similar way, *loudness* is the perception of *intensity* or *amplitude* of sound. *Loudness* is mainly dependent on amplitude, and this relationship is approximately logarithmic, which means equal ratios of amplitude are more-or-less perceived as equal increments of loudness. We are very sensitive to small ratios of frequency, but we are not very sensitive to small ratios of amplitude. To double loudness you need about a 10-fold increase in intensity. We are sensitive to about 1dB of amplitude ratio. 1dB is a change of about 12%.

The fact that we are not so sensitive to amplitude changes is important to keep in mind when adjusting amplitude. If you synthesize a sound and it is too soft, try scaling it by a factor of 2. People are often tempted to make small changes, e.g. multiply by 1.1 to make something louder, but you will likely find that a 10% change is not even audible.

Loudness also depends on frequency because we are not so sensitive to very high and very low frequencies. The Fletcher-Munson Curve, shown in blue in Figure 7, plots contours of equal loudness over a range of frequencies. These contours are lowest around 4 kHz where we are most sensitive. The curve is low here because a low amplitude at 4 kHz sounds as loud as higher amplitudes at other frequencies. As we move to even higher frequencies over on the right, we become less sensitive once again. The curves trace the combinations of frequency and amplitude that sound equally loud. If you sweep a sinusoid from low frequency to high frequency at equal amplitude, you will hear that the sound appears to get louder until you hit around 4 kHz, and then the sound begins to get quieter. Depending on how loud it is, you might stop hearing it at some point before you hit 20 kHz, even if you can hear loud sounds at 20 kHz.

The red lines in Figure 7 show an update to the Fletcher-Munson Curve based on new measurements of typical human hearing. It should be noted that these curves are based on sinusoids. In most music, low pitches actually consist of a mixture of harmonics, some of which can have rather high frequencies. So even if you cannot hear the fundamental because it is too low or too quiet, you might hear the pitch, which is implied by the upper partials (to which you are more sensitive). This is also why, on your laptop, you can hear a piano tone at the pitch C_2 (below the bass clef, with a fundamental frequency of about 65 Hz), even though your laptop speaker is barely able to produce any output at 65 Hz. You might try this SAL command, which plays a piano tone followed by a sine tone:

```
play seq(note(pitch: c2), osc(c2))
```

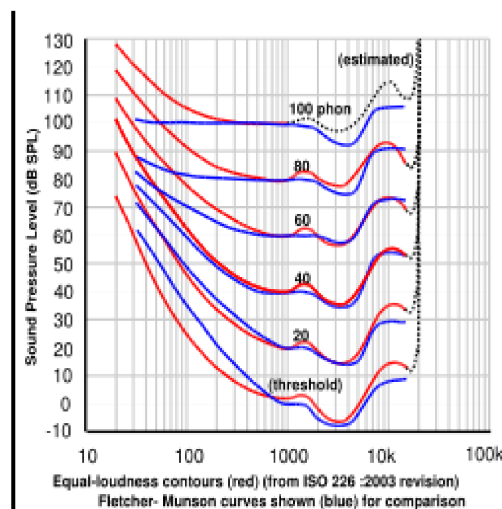

Figure 7: Fletcher-Munson Curve.

2.3 Localization

Localization is the ability to perceive direction of the source of a sound and perhaps the distance of that sound. We have multiple cues for localization, including relative amplitude, phase or timing, and spectral effects of the pinnae (outer ears). Starting with amplitude, if we hear something louder in our right ear than our left ear, then we will perceive that the sound must be coming from the right, all other things be equal.

We also use relative phase or timing for localization: if we hear something arrive earlier in our right ear than our left ear, then we will perceive that sound as coming from the right.

The third cue is produced by our pinnae or outer ears. Sound reflects off the pinnae, which are very irregularly shaped. These reflections cause some cancellation or reinforcement at particular wavelengths, depending on the direction of the sound source and the orientation of our head. Even though we do not know the exact spectrum of the source sound, our amazing brain is able to disentangle all this information and compute something about localization. This is especially important for the perception of elevation, i.e. is the sound coming from ahead or above? In either case, the distance to our ears is the same, so whether the source is ahead or above, there should be no difference in amplitude or timing. The only difference is in spectral changes due to our outer ears and reflections from our shoulders.

All of these effects or cues can be described in terms of filters. Taken together, these effects are sometimes called the HRTF, or *Head-Related Transfer Function*. (A "transfer function" describes the change in the spectrum from source to destination.) You might have seen some artificial localization systems, including video games and virtual reality application systems based on HRTF. The idea is, for each source to be placed in a virtual 3D space, compute an appropriate HRTF and apply that to the source sound. There is a different HRTF for the left ear and right ear, and typically the resulting stereo signal is presented through headphones. Ideally, the headphones are tracked so that the HRTFs can be recomputed as the head turns and the angles to virtual sources change.

Environmental cues are also important for localization. If sounds reflect off walls, then you get a sense of being in a closed space, how far away the walls are, and what the walls are made of. Reverberation and ratio of reverberation to direct sound are important for distance estimation, especially for computer music if you want to create the effect of a sound source fading off into the distance. Instead of just turning down the amplitude of the sound, the direct or dry sound should diminish faster than the reverberation sound to give the impression of greater distance.

Finally, knowledge of the sound source, including vision, recognition etc. is very important to localization. For example, merely placing a silent loudspeaker in front of a listener wearing headphones can cause experimental subjects to localize sound at the loudspeaker, ignoring whatever cues are present in the actual sound!

2.4 More Acoustics

The speed of sound is about 1 ft/ms. Sound travels at different speeds at different altitudes, different temperatures and different humidities, so it is probably more useful to have a rough idea of the speed of sound than to memorize a precise number. (But for the record, the speed of sound is 343 m/s in dry air at 20° C.) Compared to the speed of light, sound is painfully slow. For example, if you stand in front of a wall and clap, you can hear that the sound reflects from the wall surfaces, and you can perceive the time for the clap to travel to the wall and reflect back. Our auditory system merges multiple reflections when they are close in time (usually up to about 40 ms), so you do not perceive echoes until you stand back 20 feet or so from the reflecting wall.

In most listening environments, we do not get just one reflection, called an echo. Instead, we get a diffuse superposition of millions of echos as sound scatters and bounces between many surfaces in a room. We call this *reverberation*. In addition to reflection, sound refracts (bends around objects). Wavelengths vary from 50 feet to a fraction of an inch, and diffraction is more pronounced at lower frequencies, allowing sound to bend around objects.

Linearity is a very important concept for understanding acoustics. Let's think about the transmission of sounds from the source of sound to the listener. We can think of the whole process as a function F shown in the equations below, where $y = F(x)$ means that source x is transformed by the room into y at the ear of the listener. *Linearity* means that if we increase the sound at the source, we will get a proportional increase at the listening side of channel F. Thus, $F(ax) = aF(x)$. The other property, sometimes called the *superposition*, is that if we have two sources: pressure signals x_1 and x_2, and play them at the same time, then the effect on the listener will be the sum of individual

effect from x_1 and x_2:

$$F(ax) = aF(x), \quad F(x_1 + x_2) = F(x_1) + F(x_2)$$

Why does linearity matter? First, air, rooms, performance spaces are very linear. Also, many of processes we used on sounds such as filters are designed to be linear. Linearity means that if there are two sound sources playing at the same time, then the signal at the listening end is equivalent to what you get from one sound plus what you get from the other sound. Another interesting thing about linearity is that we can decompose a sound into sinusoids or components (i.e. compute the Fourier transform). If we know what a linear system does to each one of those component frequencies, then by the superposition principle, we know what the system does to any sound, because we can: break it up into component frequencies, compute the transfer function at each of the frequencies and sum the results together. That is one way of looking at what a filter does. A filter associates different scale factors with each frequency, and because filters are linear, they weight or delay frequencies differently but independently. If we add two sounds and put them through the filter, the result is equivalent to putting the two sounds through the filter independently and summing the results.

2.5 Summary of Perception

Now we summarize this discussion of acoustics and perception. Acoustics refers to the physics of sound, while perception concerns our sense of sound. As summarized in Figure 8, pitch is the perception of (mainly) frequency, loudness is the perception (mainly) of amplitude or intensity. Our perception of pitch and loudness is roughly logarithmic with frequency and amplitude. This relationship is not exact, and we saw in the Fletcher-Munson Curve that we are more sensitive to some frequencies than others. Generally, everything else we perceive is referred to as timbre, and you can think of timbre as (mainly) the perception of spectral shape. In addition to these properties, we can localize sound in space using various cues that give us a sense of direction and distance.

Perception	Acoustic Phenomenon
Pitch	Frequency (20-20 kHz range)
Loudness	Intensity (120 dB range)
Timbre	Spectrum (and other)

Figure 8: A comparison of concepts and terms from perception (left column) to acoustics (right column).

Struck objects typically exhibit characteristic frequencies with exponential decay rates (each mode of vibration has its own frequency and decay). In contrast, driven oscillators typically exhibit almost exactly periodic signals and hence harmonic spectra.

Our discussion also covered the speed of sound (roughly 1 ft/ms), transmission of sound as equivalent to filtering and the superposition principle.

3 Effects and Reverberation in Nyquist

There are a lot of effects and processes that you can apply to sound with Nyquist. Some are described here, and you can find more in the Nyquist Reference Manual. There are also many sound processing functions you can install with the Nyquist Extension Manager (in the NyquistIDE).

3.1 Delay or Echo

In Nyquist, we do not need any special unit generator to implement delay. We can directly create delay simply by adding sounds using an offset. Recall that Nyquist `sounds` have a built-in starting time and duration which are both immutable, so applying a shift operator to a sound does not do anything. However, the `cue` behavior takes a sound

as parameter and returns a new sound that has been shifted according to the environment. So we usually combine `cue` with the shift operator `@`, and a delay expression has the form:

```
cue(sound) @ delay
```

3.2 Feedback Delay

An interesting effect is to not only produce an echo, but to add an attenuated copy of the output back into the input of the effect, producing a series of echoes that die away exponentially. There is a special unit generator in Nyquist called `feedback-delay` with three parameters: `feedback-delay(`*`sound, delay, feedback`*`)`. Figure 9 shows the echo algorithm: The input comes in and is stored in memory in a first-in-first-out (FIFO) queue; samples at the end of the buffer are recycled by adding them to the incoming sound. This is an efficient way to produce many copies of the sound that fade away. The *`delay`* parameter must be a number (in seconds). It is rounded to the nearest sample to determine the length of the delay buffer. The amount of *`feedback`* should be less than one to avoid an exponential increase in amplitude.

Figure 9: Echo algorithm in Nyquist.

Note that the duration of the output of this unit generator is equal to the duration of the input, so if the input is supposed to come to an end and then be followed by multiple echos, we need to append silence to the input source to avoid a sudden ending. The example below uses `s-rest()` to construct 10 seconds of silence, which follows the *`sound`*.

```
feedback-delay(seq(sound, s-rest(10)), delay, feedback)
```

In principle, the exponential decay of the `feedback-delay` effect never ends, so it might be prudent to use an envelope to smoothly bring the end of the signal to zero:

```
feedback-delay(seq(sound, s-rest(10)), delay, feedback) * pwlv(1, d + 9, 1, d + 10, 0)
```

, where *`d`* is the duration of *`sound`*.

3.3 Comb Filter

Consider a feedback delay with a 10 ms delay. If the input is a sinusoid with 10 ms period (100 Hz), the echoes superimpose on one another and boost the amplitude of the input. The same happens with 200 Hz, 300 Hz, 400 Hz, etc. sinusoids because they all repeat after 10 ms. Other frequencies will produce echoes that do not add constructively and are not boosted much. Thus, this feedback delay will act like a filter with resonances at multiples of a fundamental frequency, which is the reciprocal of the delay time. The frequency response of a comb filter looks like Figure 10. Longer decay times gives the comb filter sharper peaks, which means the output has a more definite pitch and longer "ring."

The code below shows how to apply a comb filter to *`sound`* in Nyquist. A comb filter emphasizes (resonates at) frequencies that are multiples of a *`hz`*. The decay time of the resonance is given by *`decay`*. The *`decay`* may be a sound or a number. In either case, it must also be positive. The resulting sound will have the start time, sample rate, etc. of `textitsound`. One limitation of `comb` is that the actual delay will be the integer number of samples nearest to 1/*`hz`*.

```
comb(sound, decay, hz)
```

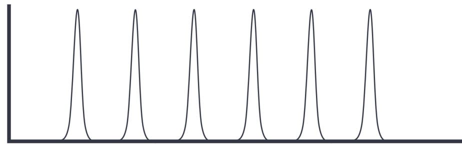

Figure 10: Filter response of a comb filter. The horizontal axis is frequency and the vertical axis is amplitude. The comb filter has resonances at multiples of some fundamental frequency.

3.4 Equalization

Equalization is generally used to adjust spectral balance. For example, we might want to boost the bass, or boost the high frequencies, or cut some objectionable frequencies in the middle range. The function `nband(`*`input, gains`*`)`, that can take an array of gains, one for each band, where the bands are evenly divided across the 20 - 20kHz range. An interesting possibility is using computed control functions to make the equalization change over time.

The Equalizer Editor in Nyquist provides a graphical equalizer interface for creating and adjusting equalizers. It has a number of sliders for different frequency bands, and you can slide them up and down and see graphically what the frequency response looks like. You can use this interface to create functions to be used in your code. Equalizers here are named `eq-0`, `eq-1`, etc., and you select the equalizer to edit using a pull-down menu. The "Set" button should be use to record changes.

The following expression in Nyquist is a fixed- or variable-parameter, second-order midrange equalization (EQ) filter:

```
eq-band(signal, hz, gain, width)
```

The *hz* parameter is the center frequency, *gain* is the boost (or cut) in dB, and width is the half-gain width in octaves. Alternatively, *hz*, *gain*, and *width* may be sounds, but they must all have the same sample rate, e.g. they should all run at the control rate or at the sample rate.

You can look up filters in the Nyquist manual for many more filters and options.

3.5 Chorus

The chorus effect is a very useful way to enrich a simple and dry sound. Essentially, the "chorus" effect is a very short, time-varying delay that is added to the original sound. Its implementation is shown in Figure 11. The original sound passes through the top line, while a copy of the sound with some attenuation is added to the sound after a varying delay, which is indicated by the diagonal arrow.

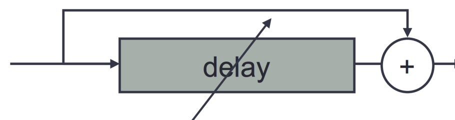

Figure 11: Chorus effect algorithm. A signal is mixed with a delayed copy of itself. The delay varies, typically by a small amount and rather slowly.

To implement chorus in Nyquist, you need to first load library `time-delay-fns`[1] and then call the `chorus` function as shown below.

```
chorus(sound, delay: delay, depth: depth, rate: rate,
   saturation: saturation, phase: phase)
```

[1] In the future, `time-delay-fns` will be an extension installed with the Extension Manager.

Here, a chorus effect is applied to *sound*. All parameters may be arrays as usual. The chorus is implemented as a variable delay modulated by a sinusoid shifted by *phase* degrees oscillating at *rate* Hz. The sinusoid is scaled by *depth*. The delayed signal is mixed with the original, and *saturation* gives the fraction of the delayed signal (from 0 to 1) in the mix. Default values are *delay* = 0.03, *depth* = 0.003, `rate` = 0.3, *saturation* = 1.0, and *phase* = 0.0 (degrees).

See also the Nyquist Reference Manual for `stereo-chorus` and `stkchorus`.

3.6 Panning

There is a detailed description of panning in the reading "Loudness Concepts & Pan Laws." In Nyquist, panning splits a monophonic signal (single channel) into stereo outputs (two channels; recall that multiple channel signals are represented using arrays of sounds), and the degree of left or right of that signal can be controlled by either a fixed parameter or a variable control envelope:

`pan(`*sound*`,` *where*`)`

The `pan` function pans *sound* (a behavior) according to *where* (another behavior or a number). *Sound* must be monophonic. The *where* parameter should range from 0 to 1, where 0 means pan completely left, and 1 means pan completely right. For intermediate values, the sound is scaled linearly between left and right.

3.7 Compression/Limiting

A Compression/Limiting effect refers to automatic gain control, which reduces the dynamic range of a signal. Do not confuse dynamics compression with data compression such as producing an MP3 file. When you have a signal that ranges from very soft or very loud, you might like to boost the soft part. Alternatively, when you have a narrow dynamic range, you can expand it to make soft sounds much quieter and louder sounds even louder.

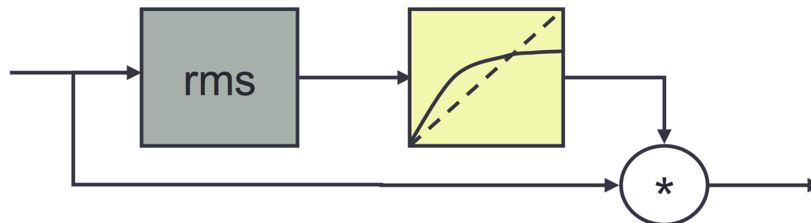

Figure 12: Compression/Limiting effect. The input signal is analyzed to obtain the RMS (average) amplitude. The amplitude is mapped to obtain a gain, and the signal is multiplied by the gain. The solid line shown in the mapping (at right) is typical, indicating that soft sounds are boosted, but as the input becomes increasingly loud, the gain is less, reducing the overall dynamic range (i.e. the variation in amplitude is reduced.)

The basic algorithm for compression is shown in Figure 12. The compressor detects the signal level with a Root Mean Square (RMS) detector[2] and uses table-lookup to determine how much gain to place on the original signal at that point. The implementation in Nyquist is provided by the Nyquist Extension named `compress`, and there are two useful functions in the extension. The first one is `compress`, which compresses *input* using *map*, a compression curve probably generated by `compress-map` [3]. Adjustments in gain have the given *rise-time* and

[2]The RMS analysis consists of squaring the signal, which converts each sample from positive or negative amplitude to a positive measure of power, then taking the mean of a set of consecutive power samples perhaps 10 to 50 ms in duration, and finally taking the square root of this"mean power" to get an amplitude.

[3]`compress-map(`*compress-ratio*`,` *compress-threshold*`,` *expand-ratio*`,` *expand-threshold*`,` *limit*`,` *transition*`,` *verbose*`)` constructs a map for the compress function. The map consists of two parts: a compression part and an expansion part. The intended use is to compress everything above compress-threshold by *compress-ratio*, and to downward expand everything below *expand-threshold* by *expand-ratio*. Thresholds are in dB and ratios are dB-per-dB. 0 dB corresponds to a peak amplitude of 1.0 or RMS amplitude of 0.7. If the input goes above 0 dB, the output can optionally be limited by setting `limit:` (a keyword parameter) to `T`. This effectively changes the compression ratio to infinity at 0 dB. If `limit:` is `nil` (the default), then the *compression-ratio* continues to apply above 0 dB.

fall-time. *lookahead* tells how far ahead to look at the signal, and is *rise-time* by default. Another function is `agc`, an automatic gain control applied to *input*. The maximum gain in dB is *range*. Peaks are attenuated to 1.0, and gain is controlled with the given *rise-time* and *fall-time*. The look-ahead time default is *rise-time*.

`compress(`*input*, *map*, *rise-time*, *fall-time*, *lookahead*`)`

`agc(`*input*, *range*, *rise-time*, *fall-time*, *lookahead*`)`

3.8 Reverse

Reverse is simply playing a sound backwards. In Nyquist, the reverse functions can either reverse a sound or a file, and both are part of the Nyquist extension named `reverse`. If you reverse a file, Nyquist reads blocks of samples from the file and reverses them, one-at-a-time. Using this method, Nyquist can reverse very long sounds without using much memory. See `s-read-reverse` in the Nyquist Reference Manual for details.

To reverse a sound, Nyquist must evaluate the whole sound in memory, which requires 4 bytes per sample plus some overhead. The function `s-reverse(`*sound*`)` reverses *sound*, which must be shorter than `*max-reverse-samples*` (currently initialized to 25 million samples). This function does sample-by-sample processing without an efficiently compiled unit generator, so do not be surprised if it calls the garbage collector a lot and runs slowly. The result starts at the starting time given by the current environment (not necessarily the starting time of *sound*). If *sound* has multiple channels, a multiple channel, reversed sound is returned.

3.9 Sample-Rate Conversion

Nyquist has high-quality interpolation to alter sample rates. There are a lot of sample rate conversions going on in Nyquist behind the scenes: Nyquist implicitly changes the sample rate of control functions such as produced by `env` and `lfo` from their default sample rate which is 1/20 of the audio sample rate. When you multiply an envelope by an audio rate signal, Nyquist *linearly* interpolates the control function. This is reasonable with most control functions because, if they are changing slowly, linear interpolation will be a good approximation of higher-quality signal reconstruction, and linear interpolation is fast. However, if you want to change the sample rate of audio, for example if you read a file with a 48 kHz sample rate and you want the rate to be 44.1 kHz, then you should use high-quality sample-rate conversion to avoid distortion and aliasing that can arise from linear interpolation.

The algorithm for high-quality sample-rate conversion in Nyquist is a digital low pass filter followed by digital reconstruction using *sinc* interpolation. There are two useful conversion functions:

`force-srate(`*srate*, *sound*`)`

returns a sound which is up- or down-sampled to *srate*. Interpolation is linear, and no pre-filtering is applied in the down-sample case, so aliasing may occur.

`resample(`*snd*, *rate*`)`

Performs high-quality interpolation to reconstruct the signal at the new sample rate. The result is scaled by 0.95 to reduce problems with clipping. (Why? Interestingly, an interpolated signal can reconstruct peaks that exceed the amplitude of the the original samples.)

Nyquist also has a variable sample-rate function:

`sound-warp(`*warp-fn*, *signal*, *wrate*`)`

applies a warp function *warp-fn* to *signal* using function composition. If the optional parameter *wrate* is omitted or NIL, linear interpolation is used. Otherwise, high-quality sample interpolation is used, and the result is scaled by 0.95 to reduce problems with clipping. Here, *warp-fn* is a mapping from score (logical) time to real time, and *signal* is a function from score time to real values. The result is a function from real time to real values at a sample rate of `*sound-srate*`. See the Nyquist Reference Manual for details about *wrate*.

To perform high-quality stretching by a fixed ratio, as opposed to a variable ratio allowed in `sound-warp`, use `scale-srate` to stretch or shrink the sound, and then `resample` to restore the original sample rate.

3.10 Sample Size Conversion (Quantization)

Nyquist allows you to simulate different sample sizes using the unit generator `quantize`:

`quantize(`*`sound`*`, `*`steps`*`)`

This unit generator quantizes *sound* as follows: *sound* is multiplied by *steps* and rounded to the nearest integer. The result is then divided by *steps*. For example, if *steps* is 127, then a signal that ranges from -1 to +1 will be quantized to 255 levels (127 less than zero, 127 greater than zero, and zero itself). This would match the quantization Nyquist performs when writing a signal to an 8-bit audio file. The *sound* may be multi-channel.

3.11 Reverberation

A reverberation effect simulates playing a sound in a room or concert hall. Typical enclosed spaces produce many reflections from walls, floor, ceiling, chairs, balconies, etc. The number of reflections increases exponentially with time due to secondary, tertiary, and additional reflections, and also because sound is following paths in all directions.

Typically, reverberation is modeled in two parts:

- Early reflections, e.g. sounds bouncing off one wall before reaching the listener, are modeled by discrete delays.
- Late reflections become very dense and diffuse and are modeled using a network of all-pass and feedback-delay filters.

Reverberation often uses a low-pass filter in the late reflection model because high frequencies are absorbed by air and room surfaces.

The rate of decay of reverberation is described by RT60, the time to decay to -60 dB relative to the peak amplitude. (-60 dB is about 1/1000 in amplitude.) Typical values of RT60 are around 1.5 to 3 s, but much longer times are easy to create digitally and can be very interesting.

In Nyquist, the `reverb` function provides a simple reverberator. You will probably want to mix the reverberated signal with some "dry" original signal, so you might like this function:

```
function reverb-mix(s, rt, wet)
  return s * (1 - wet) + reverb(s, rt) * wet
```

Nyquist also has some reverberators from the Synthesis Tool Kit: `nrev` (similar to Nyquist's reverb), `jcrev` (ported from an implementation by John Chowning), and `prcrev` (created by Perry Cook). See the Nyquist Reference Manual for details.

3.11.1 Convolution-based Reverberators

Reverberators can be seen as very big filters with long irregular impulse responses. Many modern reverberators measure the impulse response of a real room or concert hall and apply the impulse response to an input signal using convolution (recall that filtering is equivalent to multiplication in the frequency domain, and that is what convolution does).

With stereo signals, a traditional approach is to mix stereo to mono, compute reverberation, then add the mono reverberation signal to both left and right channels. In more modern reverberation implementations, convolution-based reverberators use 4 impulse responses because the input and output are stereo. There is an impulse response representing how the stage-left (left input) signal reaches the left channel or left ear (left output), the stage-left signal to the right channel or right ear, stage-right to the left channel, and stage-right to the right channel.

Nyquist has a `convolve` function to convolve two sounds. There is no library of impulse responses for reverberation, but see `http://www.openairlib.net/`, `https://www.voxengo.com/impulses/` and other sources. Convolution with different sounds (even if they are not room responses) is an interesting effect for creating new sounds.

3.12 Summary

Many audio effects are available in Nyquist. Audio effects are crucial in modern music production and offer a range of creative possibilities. Synthesized sounds can be greatly enhanced through audio effects including filters, chorus, and delay. Effects can be modulated to add additional interest.

Reverberation is the effect of millions of "echoes" caused by reflections off of walls and other surfaces when sound is created in a room or reflective space. Reverberation echos generally become denser with greater delay, and the sound of reverberation generally has an approximately exponential decay.

4 Acknowledgments

Thanks to Shuqi Dai and Sai Samarth for editing assistance.

Portions of this work are taken almost verbatim from *Music and Computers, A Theoretical and Historical Approach* (http://sites.music.columbia.edu/cmc/MusicAndComputers/) by Phil Burk, Larry Polansky, Douglas Repetto, Mary Robert, and Dan Rockmore. Other portions are taken almost verbatim from *Introduction to Computer Music: Volume One* (http://www.indiana.edu/~emusic/etext/toc.shtml) by Jeffrey Hass. I would like to thank these authors for generously sharing their work and knowledge.