Dynamics and Loudness

Topics Addressed

- Dynamics
- Decibels
- Fletcher-Munsen Curve
- Normalization
- Compression
- Limiting
- Noise Gate
- Other Topics

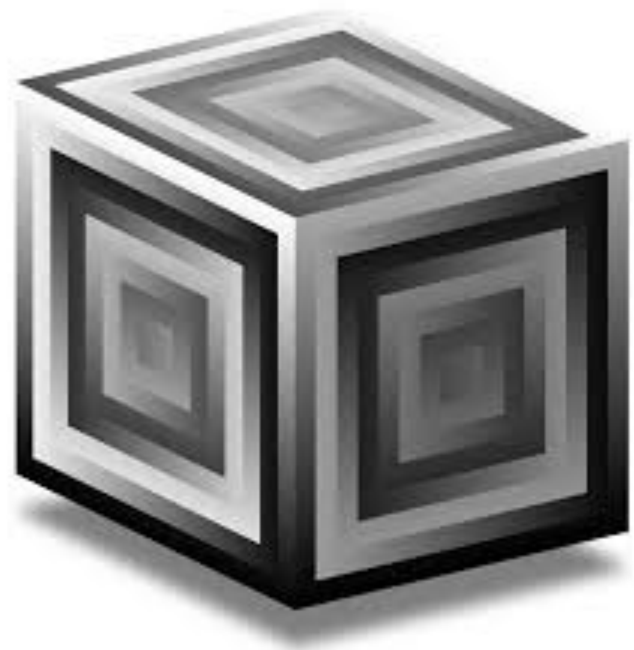

Dynamics

- What are dynamics in music?
  - The “loudness” of a particular sound relative to other sounds
- How do we measure loudness?
  - Amplitude?
- How can we control loudness in the digital world?

Loudness is Perceptual

- Human hearing is **perceptual**. There is no simple property of sound we can point to that directly correlates to our perception of loudness. Loudness is the **subjective** experience of sound pressure.
- Loudness depends on the amplitude of the frequencies
- Loudness depends upon the frequency(s) of the sound
- Loudness is roughly logarithmic
- Loudness depends upon duration
- Objective Measurement: sound pressure level (SPL)
  - Loudness is our perception of the SPL

Loudness depends upon amplitude

- Loudness depends upon the amplitude of the sound pressure.
- Instinctively, this makes sense. A sine wave with higher amplitude will sound louder than a sine wave of the same frequency at lower amplitude.

Loudness depends upon frequency

- Does loudness depend upon frequency?
- Experiment: take a sine wave and sweep from 20Hz to 20,000Hz. Does the loudness of the sound change over time or is it consistent?

Equal-Loudness Contours

- Studies first conducted by Fletcher and Munson in the 1930s
- Loudest sounds are roughly between 2-5kHz
- Phon is a logarithmic unit of loudness.
- Louder sounds change less across the frequency spectrum.

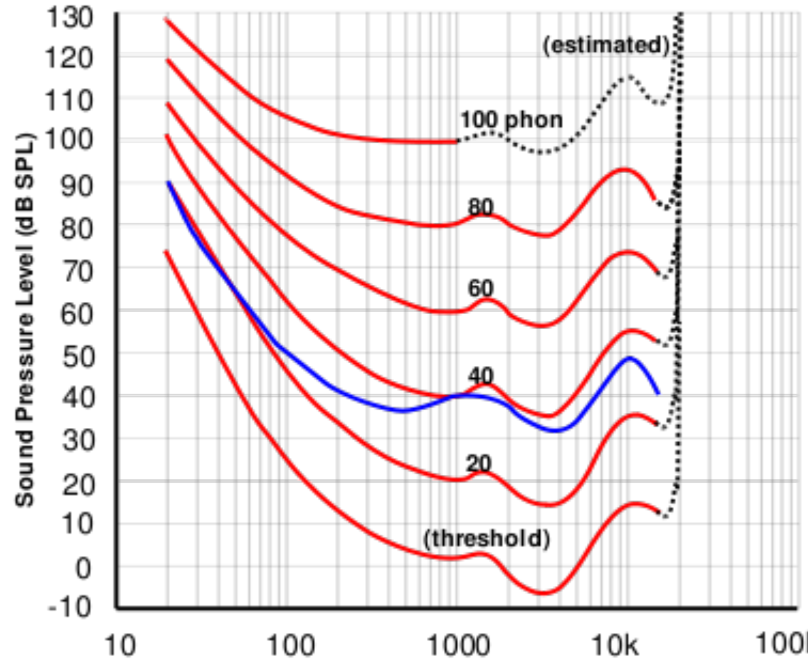

Equal-loudness contours (red) (from ISO 226:2003 revision)
Original ISO standard shown (blue) for 40-phons

Loudness is logarithmic

- Loudness is roughly logarithmic in perception.
- Doubling the amplitude of a sound increases the loudness by the same “degree” regardless of what the amplitude is.
  - For example, doubling an amplitude of 0.05 to 0.1 increases the loudness by roughly the same degree as 0.5 to 1.
  - Conversely, increasing by the amplitude by a fixed amount say 0.1 results in a different increase of loudness if the original sound has an amplitude of 0.1 vs. 0.5.
    - Question: which will have the larger **increase** in loudness, moving from an amplitude of 0.1 to 0.2 or an amplitude of 0.5 to 0.6?

Recall: Decibels

- **Decibels** are a unit of measure for measuring the ratio of amplitudes but on a logarithmic scale.
- The formula for calculating decibels is $d = 20\log_{10}(\frac{a}{a_0})$ where a is the amplitude of some signal relative to some baseline amplitude a_0.
  - In audio, a_0 is usually 1 and the measurement of our output signal will be some decibel value <= 0. This can simplify our equation to $d = 20\log_{10}(a)$
  - Alternatively, given some decibel value, we can determine the amplitude of our signal with $a = 10^{d/20}$
- 3db (shorthand for decibel) is often what we perceive as one "step" of loudness. Doubling the amplitude corresponds to about 6dB change.
- See the Wikipedia article here for a nice chart of decibels versus amplitude

Loudness depends on duration

- Loudness depends on duration
- A signal of fixed level will sound softer than if the duration is shorter.
- The threshold is true up to about 0.2/0.3 seconds after which the loudness will remain the same.

Monitoring Amplitude

- There are two common ways for displaying the amplitude of a sound: peak and RMS
- RMS of N samples from a signal x is $\sqrt{\frac{1}{N}\sum_{n=0}^{N-1} x[n]^2}$
- An audio meter (like `s.meter`) shows the fluctuation of the amplitude by either displaying the peak sample or the RMS value across some window of samples.
- Why use peak?
  - Peak signal shows the upper bounds of an audio signal.
  - Can detect if a signal goes outside the bounds of -1 to 1 to know if the signal clips.
  - SuperCollider's meter shows peak.
- Why use RMS?
  - RMS is closer to how we perceive loudness. Gets a better sense of the sustained loudness of a sound -> loudness depends upon duration

More Psychoacoustics Related to Loudness

- Masking – louder sounds often render softer sounds inaudible
- Complex sounds complicate our perception of loudness
- Predictability: unpredictable sounds are louder than predictable sounds
- Fatigue: sustained duration of loud sounds raises our threshold for hearing
  - Think of your hearing after walking out of a loud club or concert
- More: https://ccrma.stanford.edu/CCRMA/Courses/152/perceptual.html

Normalization

- Digital audio resides between -1 and 1. How can we maximize the loudness?
- Normalization: multiply all samples by a factor to expand range to maximum level
  - Algorithm:
    1. Identify maximum absolute sample across all samples
    2. Compute factor that multiplies maximum absolute sample to equal 1 (i.e., dB of 0)
        I. Note that you can normalize to other amplitudes other than 1
    3. Apply factor to every sample

Original Samples

Sample	0	1	2	3
Amp	0.1	0.5	-0.2	0.2

Normalized Samples

Sample	0	1	2	3
Amp	0.2	1.0	-0.4	0.4

Normalization Exercise

- Given the samples below, state what the output samples would be if normalized to an amplitude range of +/- 0.5

Sample	0	1	2	3
Amp	0.05	-0.1	-0.01	0.08

Dynamic Range Compression

- The idea of compression is to shrink the dynamic range of sound, generally by reducing the amplitude of louder moments.
  - Map a larger dynamic range to a smaller one
- If the amplitude passes some defined threshold in amplitude, then the signal is attenuated; otherwise the signal is passed unaltered.
- The attenuation is usually measured as a ratio where 1:1 is unaltered and ∞:1 is hard clipping at the threshold

A hard-knee compressor

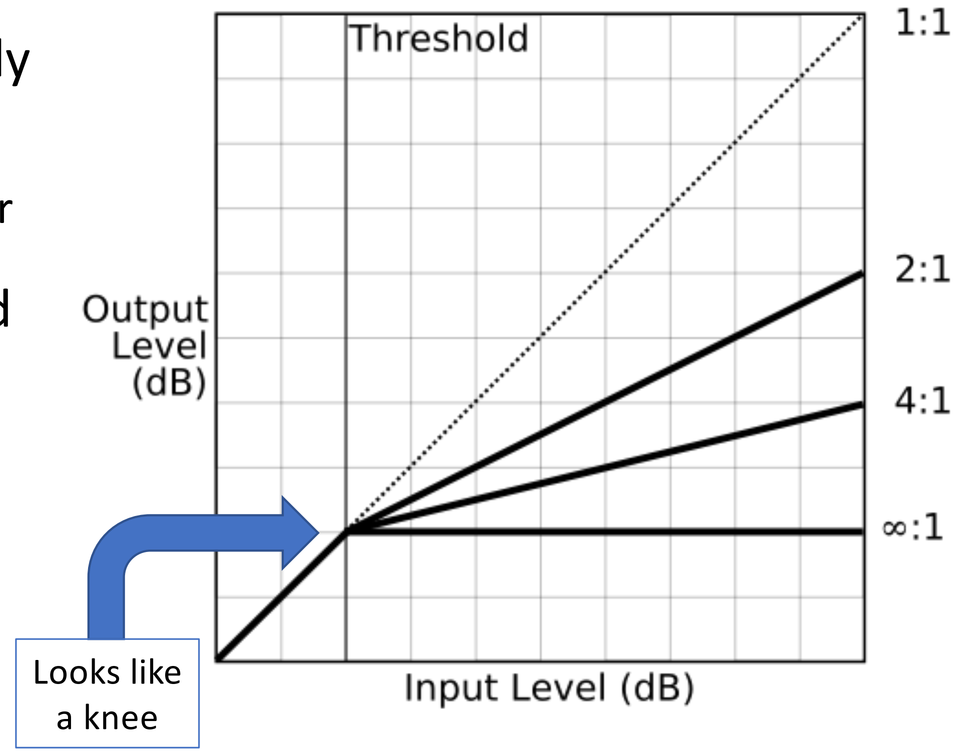

Dynamic Range Compression

Basic Algorithm:

1) Detect if level of signal is above threshold over a window of time
    a) Peak Detection
    b) RMS Detection
2) Convert signal to dB and compute gain reduction in dB using ratio
    a) Hard Knee vs. Soft Knee impacts ratio
3) Optionally apply smoothing
    a) Most compressors come with an attack/release time to transition
4) Convert gain back to amplitude and multiply by original signal to shrink dynamic range

Why Compression?

- Why compress an audio signal as an audio engineer?
  - Reducing the dynamic range makes an audio file easier to balance in a mix with other instruments.
  - Can make an audio signal louder by reducing the dynamic range and then boosting the entire signal (sometimes called “gain makeup”)
- Compression is used in hearing aids
- Broadcasting
  - Keeps uniformity of sound across different recordings
  - Audio platforms like Spotify also use compression
- Marketing/Commercials – use to achieve near-maximum loudness without achieving distortion/audio artifacts

Limiting

- Limiting is a special type of compression with a high compression ratio (10:1 or greater) and generally short attack/release time
- Limiting tends to come at the end of the signal chain to boost the overall signal but not exceed some predefined limit
- Limiting can be used in live sound situations to ensure sudden bursts of sound do not exceed some predefined limit

Recorded Audio

- Nearly all recorded audio passes through some filtering and compression
  - Filtering balances the different tracks by emphasizing different frequency ranges for each instrument
  - Compression creates a consistent loudness for each track allowing audio engineers to balance appropriately the dynamics of each instrument in the mix

The Loudness War

- An increasing trend over the decades since recorded music took off to increase the loudness
- Digital format puts a finite cap on the overall loudness of a sound (i.e., audio samples must be within range of -1 to 1)
- Compression and limiting is the way to increase the loudness of a track within this finite bound
- Interesting listen/article: https://www.npr.org/2009/12/31/122114058/the-loudness-wars-why-music-sounds-worse?sc=nl&cc=mn-20100102
- YouTube video to demonstrate: https://www.youtube.com/watch?v=3Gmex_4hreQ

1980 Super Trouper LP (POLS 322)

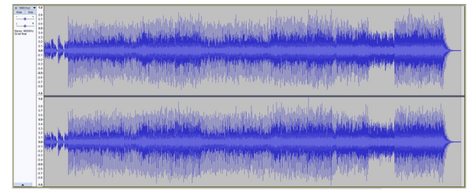

2001 Remaster (549 956-2)

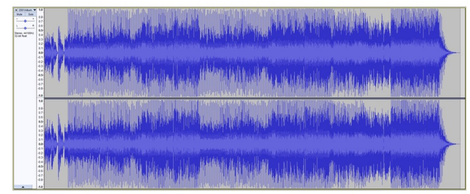

2005 TCSR Box Set (987 232-7)

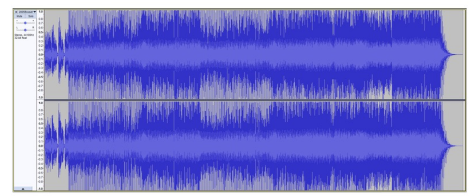

2011 Deluxe Edition (060252746446)

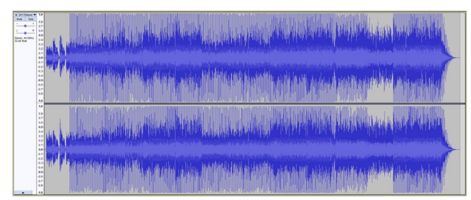

Noise Gate

- A noise gate attenuates a signal if it falls below a threshold. A common application is guitar noise/hiss from an amplifier when the guitar is not being played.
- Parameters of a noise gate:
  - Threshold – the lower bound of the signal strength that should **not** be attenuated
  - Attack – the time for the signal to ramp up to normal amplitude (i.e., opening the gate)
  - Hold – the time for signal to stay at normal volume once it initially falls below the threshold
  - Release – the time for the signal to ramp down after the hold period

Noise Gate

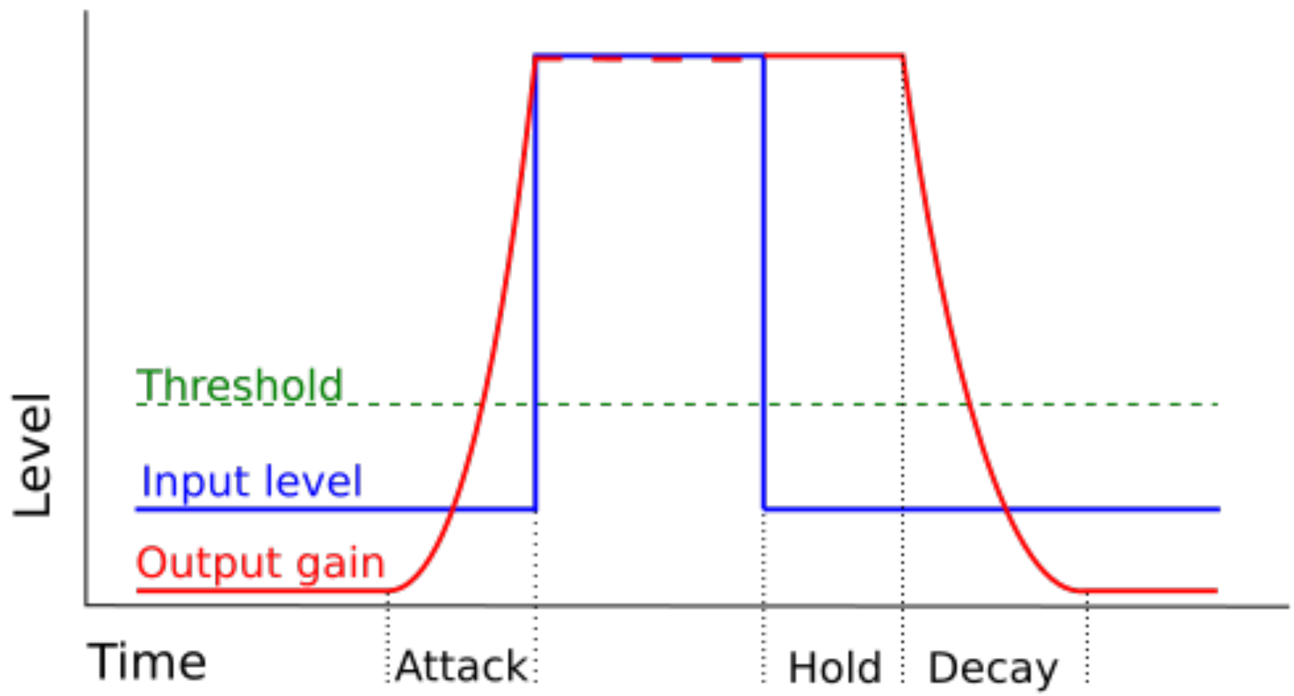
Threshold
Input level
Output gain
Level
Time
Attack
Hold
Decay

Other Topics

- Expander – expand the dynamic range – opposite of compression
- Ducking – temporarily reducing signal amplitude
  - Think reducing background music during podcasts/voiceover when speech starts
- Envelopers – essentially a compressor where you can shape the attack and release portions of the compressed sound
- De-essers – reduce volume of harsh “ess” sounds in vocals