21M.380 · Music and Technology
Recording Techniques & Audio Production

Basic sound editing techniques

Session 6 · Monday, September 26, 2016

1 Student presentation (PA1)

- [redacted]

2 Announcement: I want *you* for schlepping

- Volunteers needed for Wed, 9/28 class meeting
- 2–3 volunteers at room [redacted], 10 minutes before start of class
- 2–3 volunteers after class (please approach me after class)

3 Quiz 1 (QZ1)

4 Digital audio basics

4.1 Audio file formats

Data compression	Coding format	Container formats
Uncompressed	PCM	`.wav`, `.aif`, `.aiff`
Lossless (reversible)	FLAC ALAC	`.flac` `.m4a`

Lossy (irreversible)	MPEG layer III
	AAC
	Vorbis
	Opus
		`.mp3`
		`.m4a`, `.m4b`, `.aac`
		`.ogg`
		`.opus`

Table 1. Audio coding and container formats

Rules of 👍:

- Uncompressed for editing and production (recommended: `.wav`)
- Lossless compressed for archiving large projects (recommended: `.flac`)
- Lossy compression *only* for end-user distribution (recommended: `.ogg`)
- Avoid re-encoding in lossy compressed formats!

Value	Unit	Refers to	Application
44.1	kHz	Sample rate	Audio CD
48 000	Hz		Digital audio tape (DAT)
96	kHz		SACD, production
192	kHz		Production
16	bit	Bit depth	Audio CD
24	bit		SACD, production
32	bit		DAWs
64	bit		DAWs
128	$\text{kbit}\,\text{s}^{-1}$	Bit rate	Common `.mp3` bit rate
192	$\text{kbit}\,\text{s}^{-1}$		Common `.mp3` bit rate
256	$\text{kbit}\,\text{s}^{-1}$		High-quality `.mp3`

Table 2. Magic numbers in digital audio

4.2 Sample rate & bit depth

- Sample rate = temporal resolution (x axis; determines max. frequency)
- Bit depth = amplitude resolution (y axis; determines dynamic range)
- General recommendation for this class: 44.1 kHz, 24 bit
  - 24 bit probably more important than 96 kHz
  - Stick to one sample rate & bit depth for given project (avoid unnecessary conversions)
- Bit rate … describes data throughput of compressed audio formats
  - Unit indicates product of bit depth (kbit) & sample rate (s^{-1} = Hz)
  - In reality not quite as simple (since data is compressed)
  - Modern codecs often use *Variable bit rate* (VBR) to minimize file size

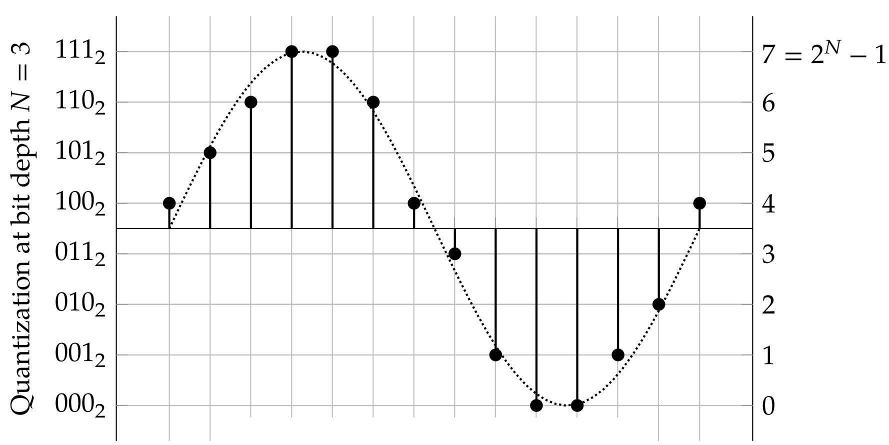

Figure 1. Sample rate and bit depth in digital audio

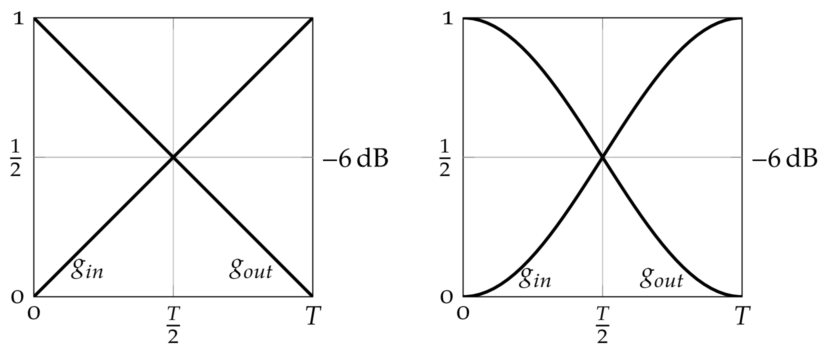

(A) Linear

(B) Half cosine

Figure 2. Symmetrical constant-gain crossfades ($g_{in} + g_{out} = 1$)

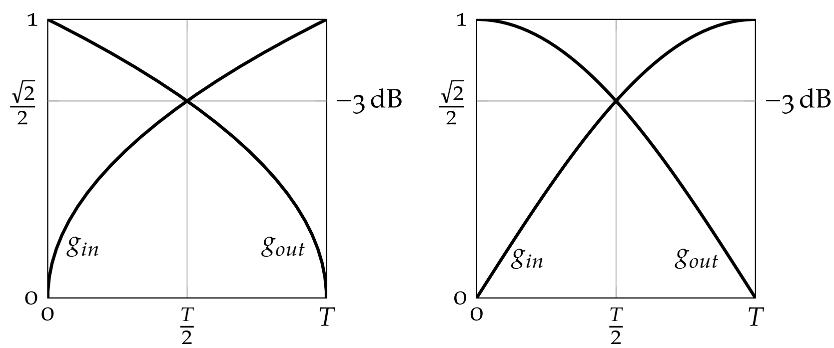

(A) Square-root

(B) Quarter cosine

Figure 3. Symmetrical constant-gain crossfades ($g_{in} + g_{out} = 1$)

5 Essential sound editing principles

5.1 Avoiding clicks through crossfades

- Rapid changes in amplitude results in audible *clicks*. Why?
- Likely to happen whenever you splice together two audio snippets!
- Use *crossfades* to prevent clicks
- Different crossfade shapes available. Rules of 👍:
  - Use constant-gain ($g_{in} + g_{out} = 1$) for strongly correlated materials
  - Use constant-power ($g_{in}^2 + g_{out}^2 = 1$) for uncorrelated materials
- Crossfades can also be asymmetrical (fade-out time of first sound ≠ fade-in time of second)

5.2 Setting appropriate output levels & avoiding clipping

- Hard border of digital audio is 0 dB_{FS} ("dB full-scale")
  - Violating that border results in *clipping* of waveform

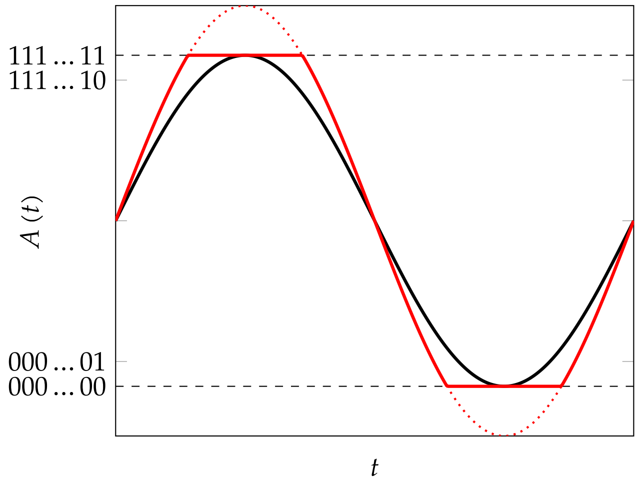

Figure 4. Full-scale (black) and clipped (red) digital signal

– Audible result: Harmonic distortion (i.e., additional high-frequency components). Why distortion? Why harmonic?

• But don't be tempted to record or render at low levels 'to be safe' and amplify later. Why?

• Rules of 👍 for recording, editing & mixing:

– Always aim at getting good levels *as early as possible* in the process

– Make levels 'as high as possible, but not higher' to optimize S/N ratio

• Rules of 👍 for mastering:

– Don't render *all* the way to 0 $\mathrm{dB_{FS}}$ (inter-sample peaks; also depends on DAC)

– Normalize to a *loudness target* (discussed below) suitable for target medium

• Recommendation for ED assignments: Render to −3 dB to −1 dB peak level

5.3 Normalization

• Motivations:

– Exploit available headroom of low-level signals

– Match perceived loudness of different audio files (album; playlist)

– Win the *loudness war* (Katz 2014a) ☹

• Peak or RMS *normalization* is a two-step process:

1. Find maximum (peak or RMS) level in signal to be normalized
2. Amplify (or attenuate) signal to a new pre-defined target maximum

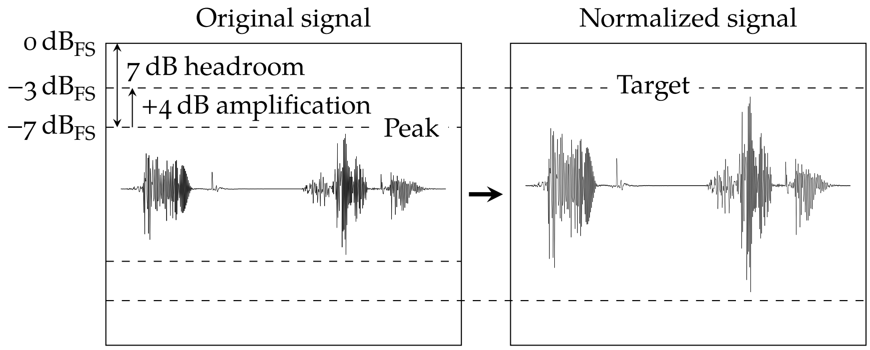

Figure 5. Principle of peak normalization

- Example:
  - Maximum peak level is −7 dB
  - We want to normalize to −3 dB peak level
  - Hence entire signal has to be amplified by +4 dB
- Properties:
  - Normalization can be automated
  - Inherently non-realtime!
  - Only changes overall *level* (but neither spectrum nor dynamics!)
- Modern productions use *loudness normalization* (Katz 2014a,b,c)
  - Neither peak nor rms level say much about perceived loudness ☹
  - Hence nobody knew how their production's loudness would compare to others
  - So everybody tried to make theirs as loud as possible → *loudness war*
  - Will be discussed in more depth in future lecture on mastering
- Lufs ("loudness unit with regards to full scale")
  - Models perceived loudness for purpose of loudness normalization
  - Target recommendations available, e.g., -23 lufs for tv broadcasts
  - Details: Itu-r bs.1770-4 (Itu 2015); ebu r128 (Camerer 2011)
- When to normalize?
  - Peak/rms: When you have a very low-level audio file where you reach the top of the fader in the mix
  - Loudness normalization: Before rendering your mix (or during mastering)
  - But no need to normalize each and every individual audio snippet!

- How to normalize?
  - As a destructive offline process (e.g., `sox in.wav out.wav norm -3`, or Effect » Normalize menu in Audacity)
  - Manually. Project renders to meager peak level of −13 dB? Raise master fader by 12 dB and re-render to −1 dB peak level!
  - Automatically. Many DAWs provide functions to automatically normalize individual audio snippets or master fader

5.4 Basic mixing: Panning & level balance

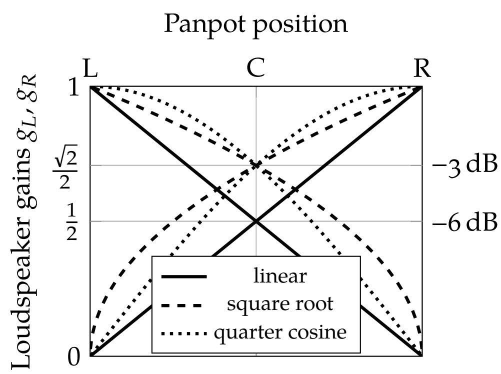

Figure 6. Stereo panpot laws

- Basic act of *mixing*: Setting level balance and stereo balance:
- Faders adjust level at which input signal is routed to stereo mix
- *Panpot* determines position of a mono *phantom source* in stereo image
  - *Panning* refers to positioning of sound between L & R speakers
  - Reciprocally adjusts gains of 2 amps for L & R speakers (*amplitude panning*, by contrast to *time delay panning*)
  - Implemented on analog mixers by means of a *potentiometer*, hence the term *panpot*
  - Different panpot laws determine precise gain functions (usually *equal power panning*)
- Panpot vs. balance control:
  - Panpot: Mono input signal (entire input signal will be heard regardless of panpot position)
  - Balance control: Stereo input signal (parts of input might become inaudible for hard L/R panning!)

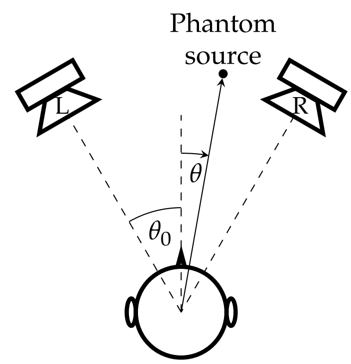

Figure 7. Stereo loudspeaker setup

6 Basic sound editing in Reaper

6.1 Licensing information

- $60 for personal use, educational institutions, small businesses
- Free 2-month trial

6.2 Download

- `http://reaper.fm/download.php`
- On Linux, download Windows installation file for installation via Wine
- Do I have a 32 or 64 bit os? Probably 64 bit, but:
  - macOS: Apple > About This Mac (top left desktop corner)
  - Windows 8 and 10: Win + X and System > System type
  - Linux: Type `arch` in command line (`i686` = 32 bit, `x86_64` = 64 bit)
- Reaper user guide (actually useful!): `http://reaper.fm/userguide.php`

6.3 Installation

- macOS: Drag from `.dmg` to 🗀 `Applications`
- Windows: Double-click `.exe`
- Linux:
  1. On 64-bit machines, enable multiarch, e.g.:
     `sudo dpkg --add-architecture i386 && sudo apt-get update`
  2. Install Wine and Mono, e.g., via Ubuntu Software Center or
     `sudo apt-get install wine winetricks mono-complete wine32`

6.4 Overview of Reaper window

- Track control panel (left)
- Transport bar (below)
- Arrange area (right)
- Mixer control panel (bottom)

6.5 Recording sound

- Add new track(s): Tracks > Insert new track
- Arm track(s) for recording:
  1. Click small red circular button on respective track
  2. Track meter should now show input signal
- Start recording: Large red circular button (transport area) or Ctrl + R
- Stop recording: Click record button again (or stop button)

6.6 Importing audio files

- Insert › Media file...
- Can also select *multiple* sound files and insert them as Separate Tracks

6.7 Playing back sound

- Start playback using ▶ button or (preferably) Space bar
- Hit Space again to stop playback and go back to starting point
- Ctrl+Space to stop playback and leave cursor in place
- R key (or use button in transport menu) to toggle looped playback
- Mute a channel by pressing M button in track control panel
- Solo a channel by pressing S button in track control panel

6.8 Mixing & panning

- Input channel faders (on track control panel *and* mixer)
- Panpot (on track control or mixer control panel)
- Left: master fader (mixed stereo output)

6.9 Arranging

- Drag & drop items on timeline or between tracks (left mouse button)
- Hold Ctrl to select multiple items with left mouse button
- Disabling grid lines:
  - Alt+G (Win)
  - ⌥+G (Mac)
  - Details: *Displaying Grid Lines* in *Reaper User Guide* (Francis 2017)

6.10 Cutting items

1. Select item(s) to cut
2. Position cursor on timeline and press
3. s key to cut

6.11 Shortening, lengthening & looping

1. Move cursor over bottom half of left or right edge of audio snippet
2. Hold left mouse button and move mouse left/right

Note: Extending item beyond its actual length will automatically loop it

6.12 Splicing, fades & crossfades

- Fade in | out:
  1. Move cursor over top half of left | right edge of audio item
  2. Left-click and drag mouse to right | left to set fade length
- Crossfades:
  - Reaper automatically creates crossfade between overlapping items
  - Double-click on crossfade area to edit in detail

6.13 Rendering a completed project to a new sound file

- Don't submit `.RPP` (Reaper project) files – they contain *no* audio data!
- Instead, render completed project to new audio file via File ⟩ Render...:
  - Select channels to render (usually `Master mix`)
  - Select render bounds (avoid mysterious periods of silence at end)
  - Decide on output file format, sample rate, bit depth, channel number
  - Watch output level meter
  - Adjust master fader, and re-render if necessary (to avoid clipping)

7 Preview ED1 assignment

References & further reading

Camerer, Florian (2011). *Practical Guidelines for Production and Implementation in Accordance with EBU R 128*. EBU Tech 3343. European Broadcasting Union. 44 pp. URL: `http://tech.ebu.ch/docs/tech/tech3343.pdf` (visited on 09/07/2013).

Francis, Geoffrey (2017). *Up and Running. A Reaper User Guide*. URL: `http://reaper.fm/userguide.php` (visited on 09/19/2017).

International Telecommunication Union (Oct. 2015). *Recommendation ITU-R BS.1770-4. Algorithms to measure audio programme loudness and true-peak audio level*. URL: `https://www.itu.int/dms_pubrec/itu-r/rec/bs/R-REC-BS.1770-4-201510-I!!PDF-E.pdf` (visited on 06/04/2017).

Katz, Bob (2014a). "The loudness revolution. The war is ending." In: *Mastering Audio. The Art and the Science*. 3rd ed. Burlington, MA: Focal Press. Chap. 17, pp. 241–56. ISBN: 978-0240818962. MIT LIBRARY: `002307049`. On course reserve at the Lewis Music Library.

— (2014b). "The loudness revolution. Loudness metering: it's time." In: *Mastering Audio. The Art and the Science*. 3rd ed. Burlington, MA: Focal Press. Chap. 18, pp. 257–61. ISBN: 978-0240818962. MIT LIBRARY: `002307049`. On course reserve at the Lewis Music Library.

— (2014c). "The loudness revolution. Calibrated monitoring." In: *Mastering Audio. The Art and the Science*. 3rd ed. Burlington, MA: Focal Press. Chap. 19, pp. 263–72. ISBN: 978-0240818962. MIT LIBRARY: `002307049`. On course reserve at the Lewis Music Library.