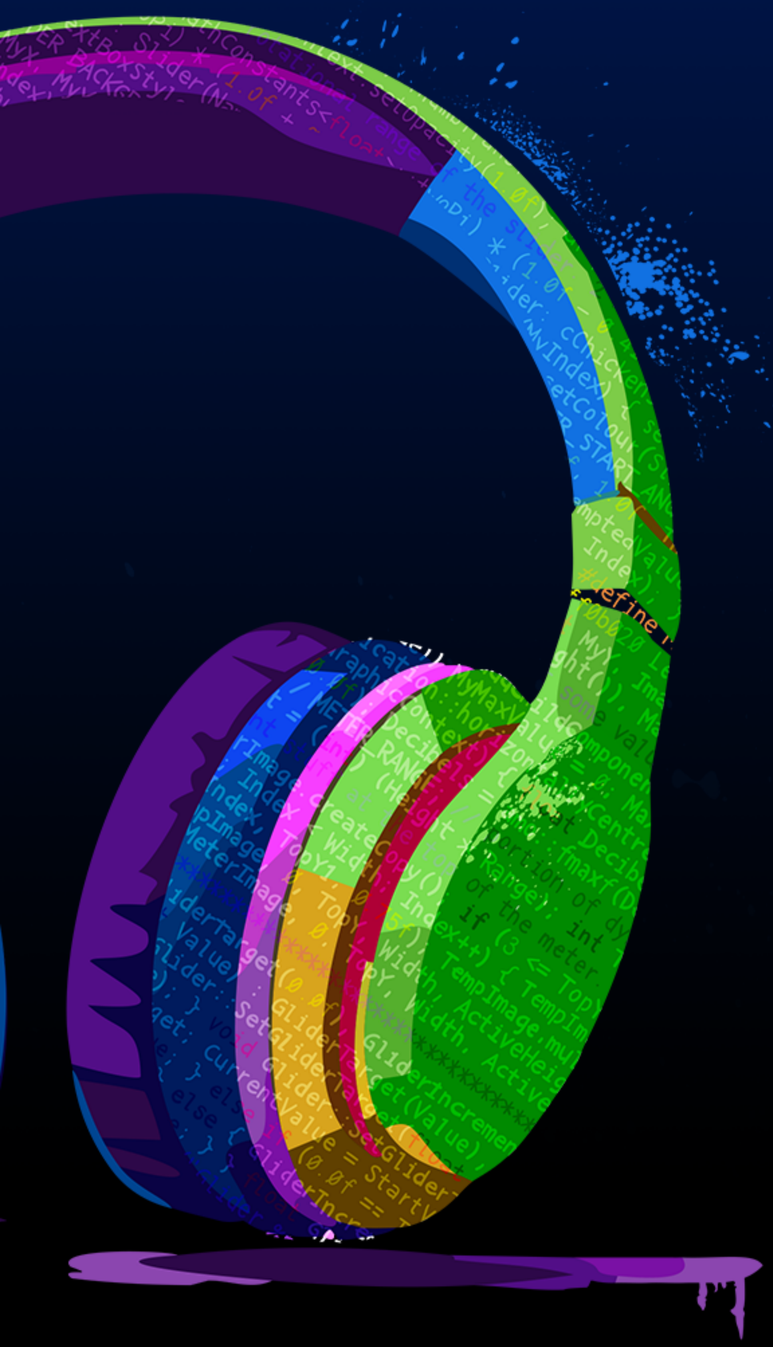

ADC24 Bristol

IN PRAISE OF LOUDNESS

SAMUEL FISCHMANN

In Praise of Loudness

Sam Fischmann

MUSIK HACK

Who is this guy?

Co-Founder of Musik Hack

20+ years software development

Longtime musician, composer, and cook

Dog person / allergic to cats

MUSIK HACK

What is this talk about?

What is loudness and how is it achieved?

How is it measured?

What are the standards related to it?

What are the myths?

MUSIK HACK

The Basics of Digital Loudness

CRASH COURSE!

MUSIK HACK

TOO LOUD

NOTDYNAMICENOUGH

tOodyna MIc

CLIPPED

MUSIK HACK

Loudness is Movement

Sound is **movement**

Faster movement is **higher pitch**

More peak to peak travel is **louder**

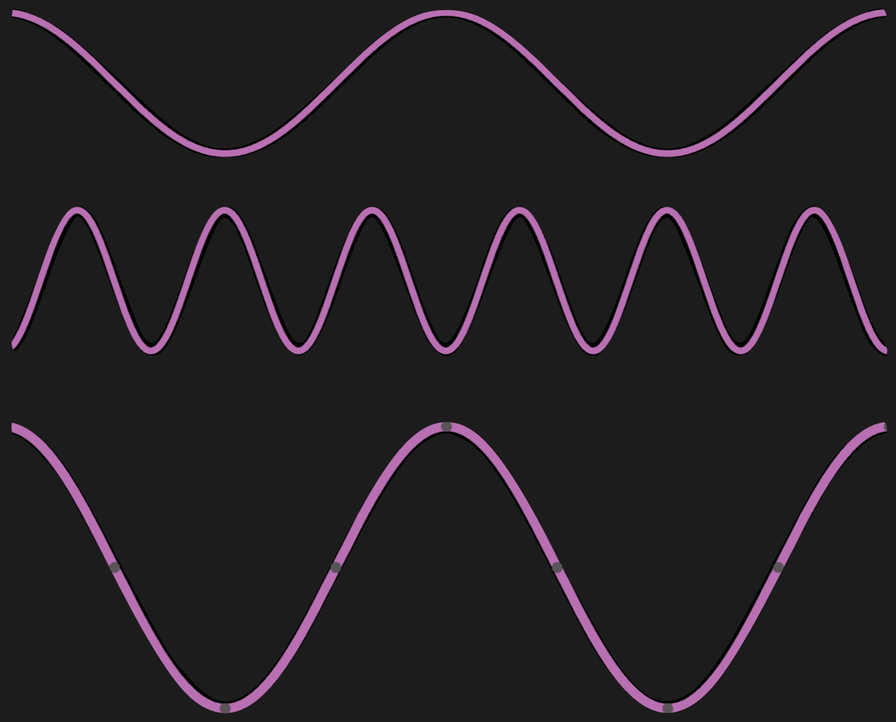

MUSIK HACK

Signal-to-Noise Ratio (SNR)

Noise can overpower signal

Signal can overpower noise

But they don’t *always* interfere

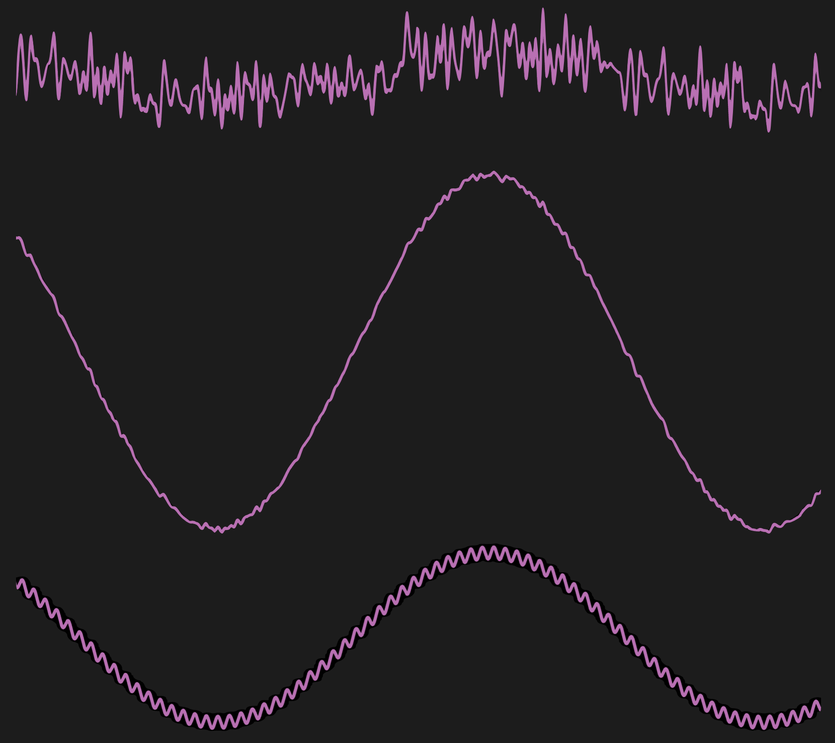

MUSIK HACK

Digital Movement and Bit Depth

DAW Range: -1.0 to 1.0 (floating point, 32- or 64-bit)

Asset Range: 16- or 24-bit integers

Higher bit depth == more precise samples

Loudness is movement, so **quiet** means **less resolution**

MUSIK HACK

Quiet Considerations: Quantization

All digital audio “snaps” to the sample grid

Diff from real to snap is **quantization noise**

This is why we mix in 32/64 bits

This is why we invented dithering!

MUSIK HACK

Loudness vs. Dynamic Range

Loudness describes one window of sound in time

Dynamic range compares at least two windows

… or, in the case of bit depth the quietest and loudest signals possible

MUSIK HACK

Loudness vs. Dynamic Range

Decibels is the measurement for loudness

Relative unit of measurement, logarithmic, 10 dB ~2x

dBFS relative to 0 as max peak to peak movement

16-bit: 96 dB, 24-bit 144 dB theoretical, ~130 max real

Dynamic range in vinyl is 50-70dB, usually closer to 60 on a decent system

MUSIK HACK

Human Concerns: Contours, Pink Noise, etc.

Different frequencies are perceived differently

Pink noise vs. white noise

Fletcher Munson

ISO 226

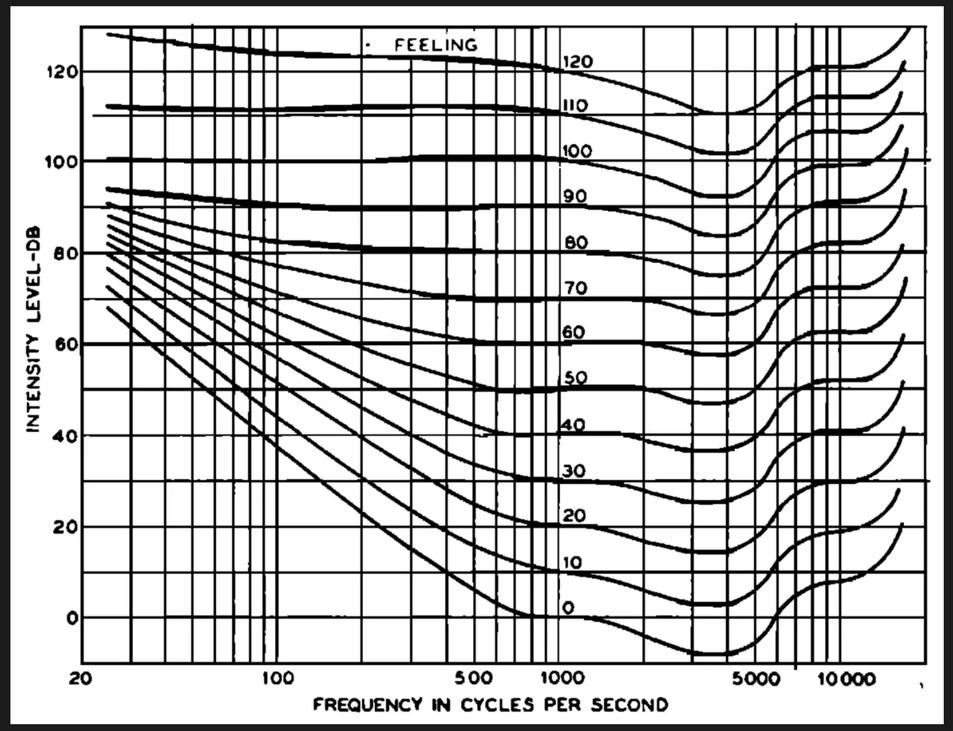

MUSIK HACK

Measurements

THE LANGUAGE OF THE “THIEF OF JOY”

MUSIK HACK

Amplitude vs. Decibel Full-Scale

Amplitude is a raw sample value, decibel is a perceptual log-scale value

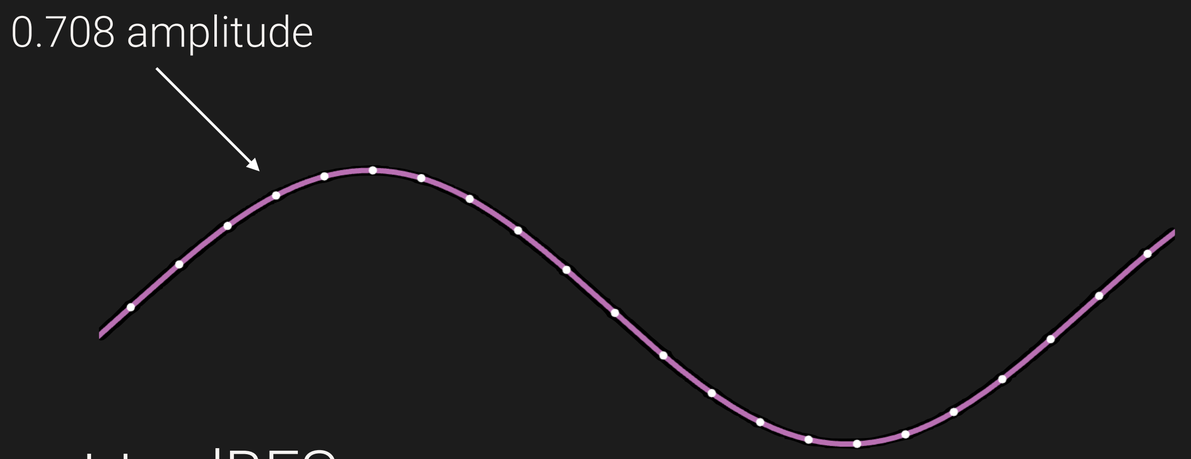

To convert to dBFS:

```
20 * log10(amplitude)
```

0.708 == -3 dBFS

MUSIK HACK

I’m going to use dB instead of dBFS from now on for brevity. But I mean dBFS.

MUSIK HACK

Loudness vs. Amplitude

1.0, or 0 dB (amplitude)

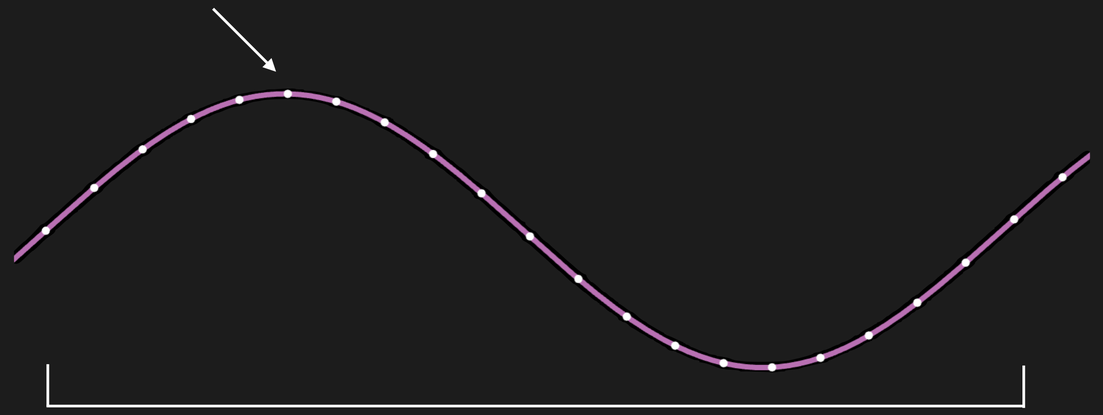

<aggregation & math>

0.708, or -3 dB (loudness)

MUSIK HACK

Peak

The immediate value of the sample farthest from 0

MUSIK HACK

Intersample Peak

The theoretical actual peak amplitude after reconstruction

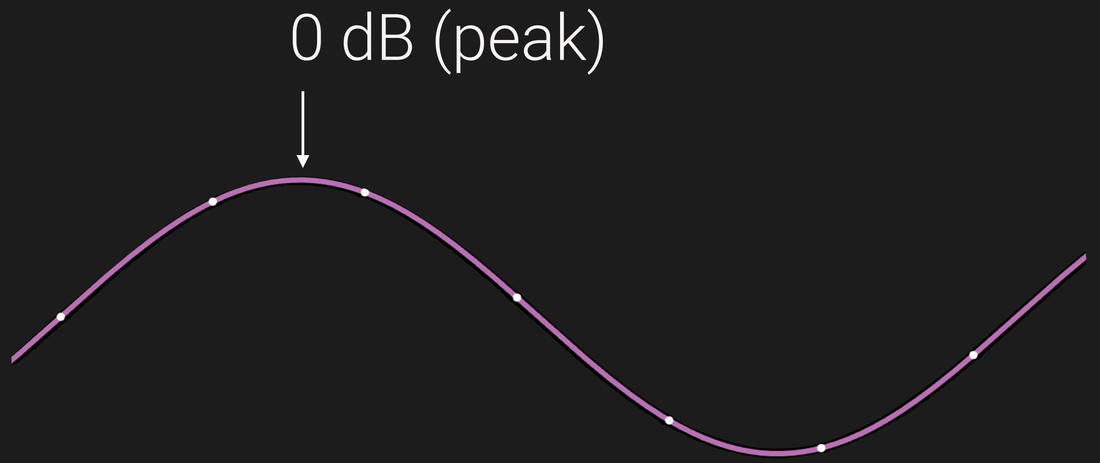

MUSIK HACK

“True” Peak

Standardized intersample peak

Defined in ITU-R BS.1770

4x oversampling w/ specified filter *or better*

I did a talk about oversampling: https://www.youtube.com/watch?v=yJrJi-4SDUQ

MUSIK HACK

RMS: Root Mean Squared

The classic loudness measurement

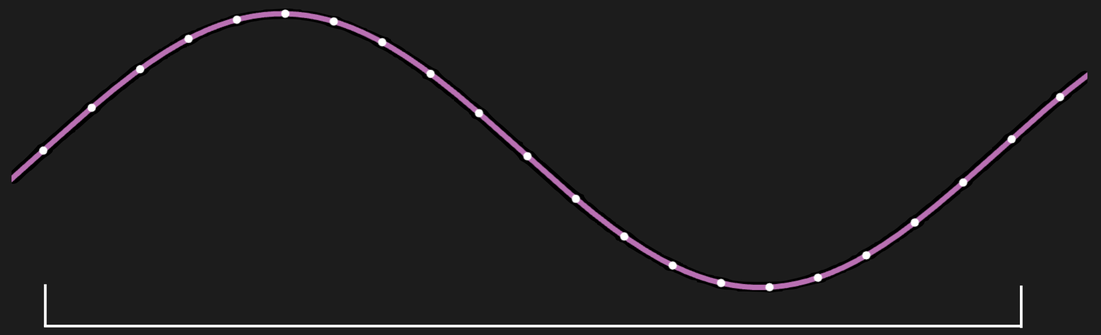

300 ms window
(VU meter)

square, mean(), sqrt()

$$\sqrt{\frac{1}{n}\sum_{i=1}^{n} x_i^2}$$

0.708, or -3 dB

MUSIK HACK

LUFS

Most popular modern loudness measurement

Perception-based (frequency loudness curves)

Ignores very quiet or silent material

Extended to multichannel loudness*

*but Atmos object-based mixing must be rendered out to 5.1 before measurement, or for bed-based, you've got some calculations to do depending on the output system, defeating the purpose of the media format to begin with: https://www.itu.int/dms_pubrec/itu-r/rec/bs/R-REC-BS.2051-3-202205-I!!PDF-E.pdf

MUSIK HACK

LUFS / LKFS

1. “K” filter: high shelf ~1.2 kHz, low cut ~100 Hz
2. Break into 75% overlapping blocks of 400 ms
3. Square, mean, convert to LUFS for block
4. Discard blocks under absolute threshold of -70
5. Avg. remaining blocks, subtract 10 for relative threshold
6. Discard all blocks under the relative threshold
7. Avg. remaining blocks for the LUFS value of the segment

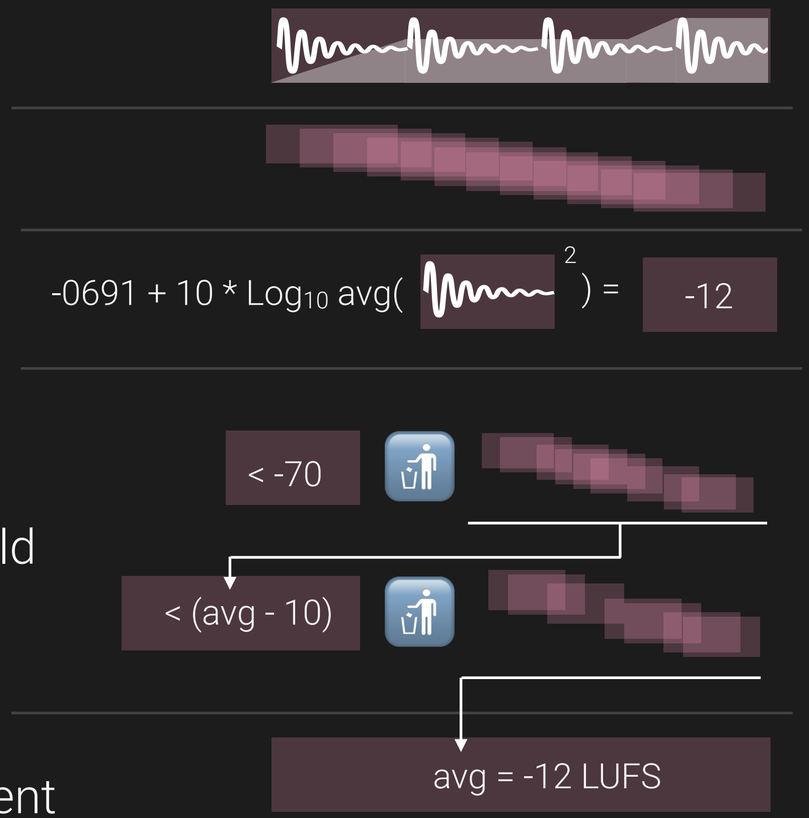

MUSIK HACK

LUFS Variants

LUFS-M / Momentary: a single 400 ms block

LUFS-S / Short-term: over 3 seconds (usually)

LUFS-I / Integrated: over a whole piece of media

MUSIK HACK

Loudness Range (LRA)

The macro-level dynamics of a signal

After discarding blocks below thresholds, find:

- The loudest block in the quietest 10th percentile
- The quietest block in the loudest 95th percentile

Calculate the LUFS difference between them!

MUSIK HACK

Peak to Loudness Ratio (PLR) & Crest Factor

The micro-level dynamics of a signal (punctuation)

PLR: True peak / Integrated LUFS

Crest: Peak / RMS (usually)

MUSIK HACK

Limiters, Clippers, Compressors

OH MY!

MUSIK HACK

Characteristics of Loudeners

Do transients get brighter or darker?

Does it add saturation, when, and how much?

Does it “pump” or “fart” when pushed?

Does it re-balance the mix?

MUSIK HACK

Saturator

Transfer function

Adds distortion

Odd-order harmonics are cleaner (3rd)*

Nonlinear within normal range/aliasing

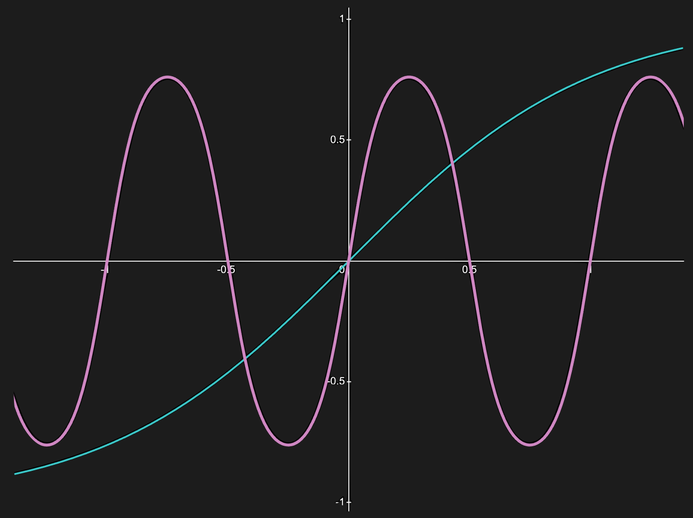

*See my talk for the DSP Online conference

MUSIK HACK

Soft Clipper

Transfer function

Similar to a saturator, hard limit

Nonlinear within normal range/aliasing

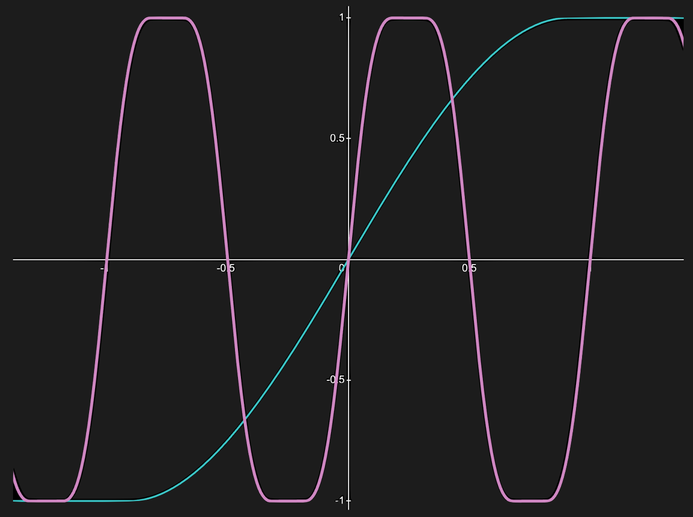

MUSIK HACK

Hard Clipper

Transfer function

Linear within allowable range (-1,1)

Sudden distortion burst on overs

Bursts can “hide” inside transients

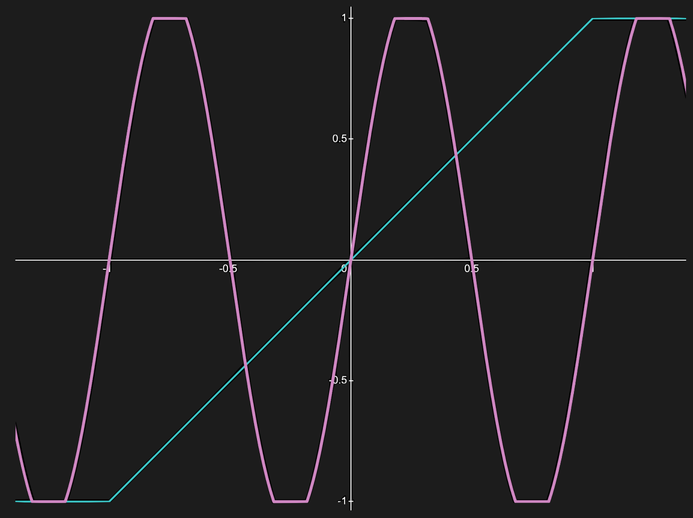

MUSIK HACK

Compressor

Linear within threshold, non-linear outside

Ballistics section gradually applies/releases gain reduction

The longer outside the threshold, the more intense

Lots of decisions with wide impacts on sound:

Log or linear domain
Peak vs RMS
Knee shape and type

feed forward or backward
ballistics smoother type
attack and release times

lookahead amount
max dB reduction
every curve tunable

MUSIK HACK

Limiter

Basically, a compressor that prevents clipping (but still flatlines)

All design decisions from compressor slide

Extensive use of lookahead and fast attack times

Release tuning is the really exciting part

MUSIK HACK

Multiband Compressor

Linkwitz Reily (or similar) crossover to split into bands

Compress each band, then rejoin

Alters the character of the mix when it gets loud

Another variation on the idea: dynamic EQ

MUSIK HACK

🔥 Unique Processes 🔥

MUSIK HACK

All this to say...

There are lots of ways to make sound louder

They do not all sound the same

Some are more transparent, some are more creative

All of them are ***TOOLS***

MUSIK HACK

About Lookahead

Limiters “duck” volume change

Zero attack either hard or soft clips (hard or soft knee)

Lookahead allows ducking ***before***, retaining some shape

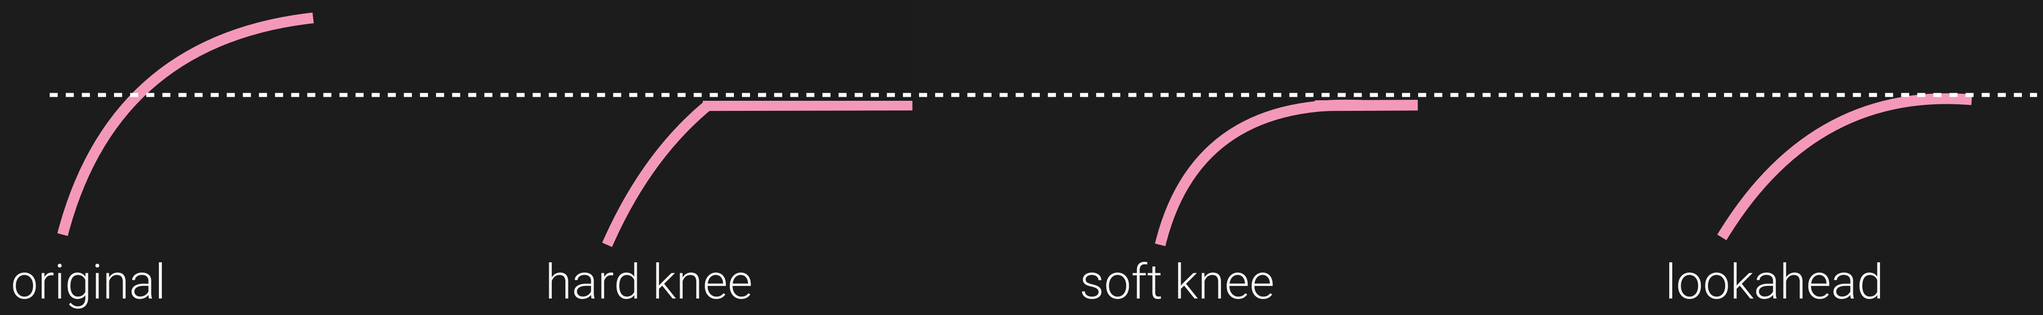

MUSIK HACK

The Loudness Wars

IF YOU CAN'T BEAT 'EM, LOOK AHEAD

MUSIK HACK

Pre-CD

Mastering/Finalizing audio for the format (cassette, vinyl)

Limitations on volume for device tolerance

MUSIK HACK

CD

Perfect replication for distribution

Peak-normalization paradigm for compatible delivery

Well-defined dynamic range and frequency characteristics

Less variability on playback (until amplification stage)

MUSIK HACK

Lookahead Limiters

Lookahead meant peaks could be ducked without poke through

Peak-normalization’s only limitation is peak below 0

Engineers start to push boundaries for radio dominance

But also: radio is already compressed, engineers want control

MUSIK HACK

Sreedhar / IEEE Spectrum / “The Future of Music”

Criticized modern music production, not able to “breathe”

Focused on a flat waveform; visual indication in micro dynamics

Acknowledges listening environment as drive for loudness

Perplexingly mentions higher-bit-depth formats as a solution?*

*”Audiophiles looking to the future for relief from overcompression see a cloudy picture. DVD-Audio and Super Audio Compact Disc (SACD) are two high-fidelity formats that were thought to be solutions to the loudness war. Both formats offer not only a greater dynamic range than CD but also higher sampling rates.” - Suhas Sreedhar, The Future of Music

MUSIK HACK

Sentiments in “The Future of Music”

“Even though you love this album, you can't listen to it anymore. You shut it off, tired, puzzled, and confused.”

“… it could be responsible for halting technological advances in sound quality for years to come.”

“Not only is all impact lost, but the constant level of the sound is fatiguing to the ear”

MUSIK HACK

Deruty / Sound on Sound

Uses LRA instead, criticizes Sreedhar for focusing on crest

Says macro dynamics has not changed

Extensive research on old tracks and new

MUSIK HACK

Ian Shepard / Bob Katz

Deruty is wrong; we are explicitly talking about micro dynamics

Suggests using PLR instead of LRA

Started “Dynamic Range Day”

They’re right about Deruty’s argument, but not necessarily right

MUSIK HACK

THEN WHO IS RIGHT?

You can’t use measurements to insist on what’s better

Just look at John Atkinson in Stereophile reviews!

Remember Sreedhar’s thought about listening environments?

Also, loudness ***is a creative decision***, like harmony or dissonance

MUSIK HACK

It’s just PLR/transients at issue? Let’s listen...

Original Sample

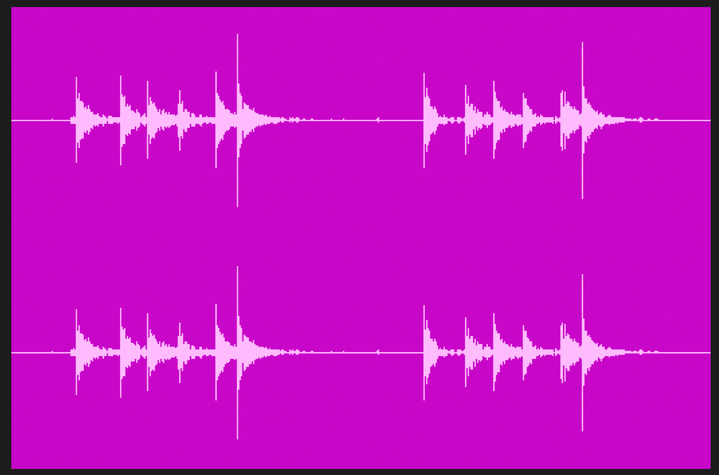

https://www.cambridge-mt.com/ms/mtk/

MUSIK HACK

It’s just PLR/transients at issue? Let’s listen...

+6 dB, gain matched

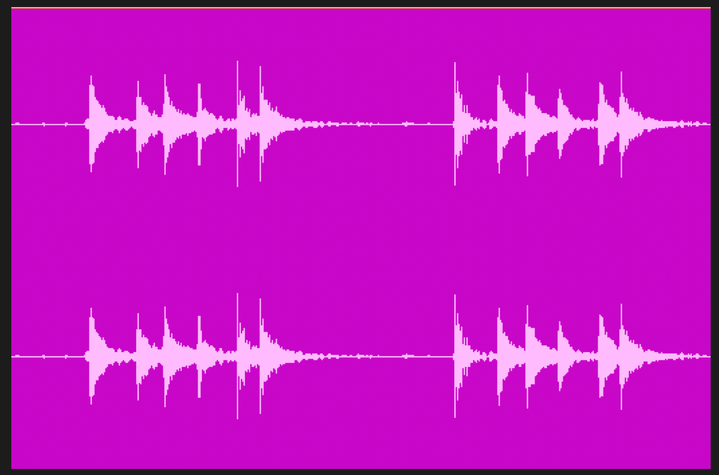

MUSIK HACK

It’s just PLR/transients at issue? Let’s listen...

+12 dB, gain matched

MUSIK HACK

It’s just PLR/transients at issue? Let’s listen...

+12 dB Gain applied

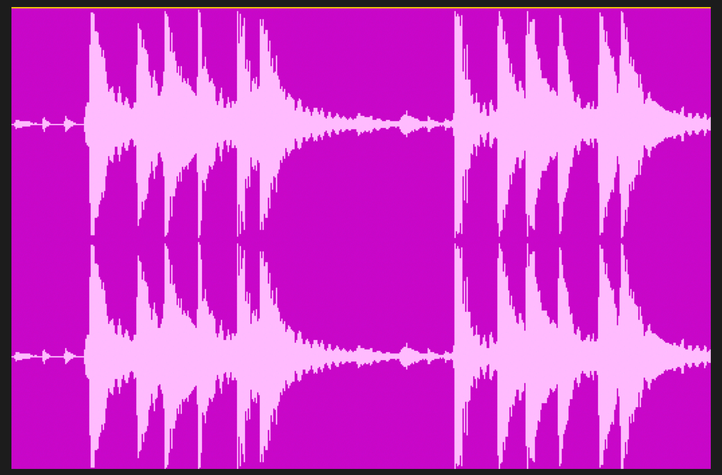

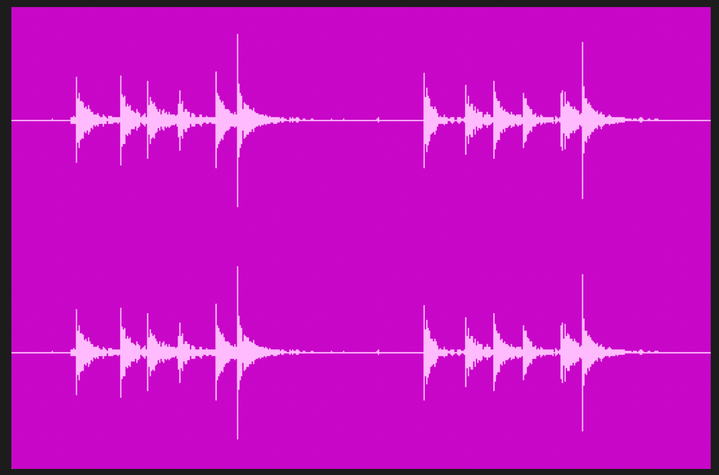

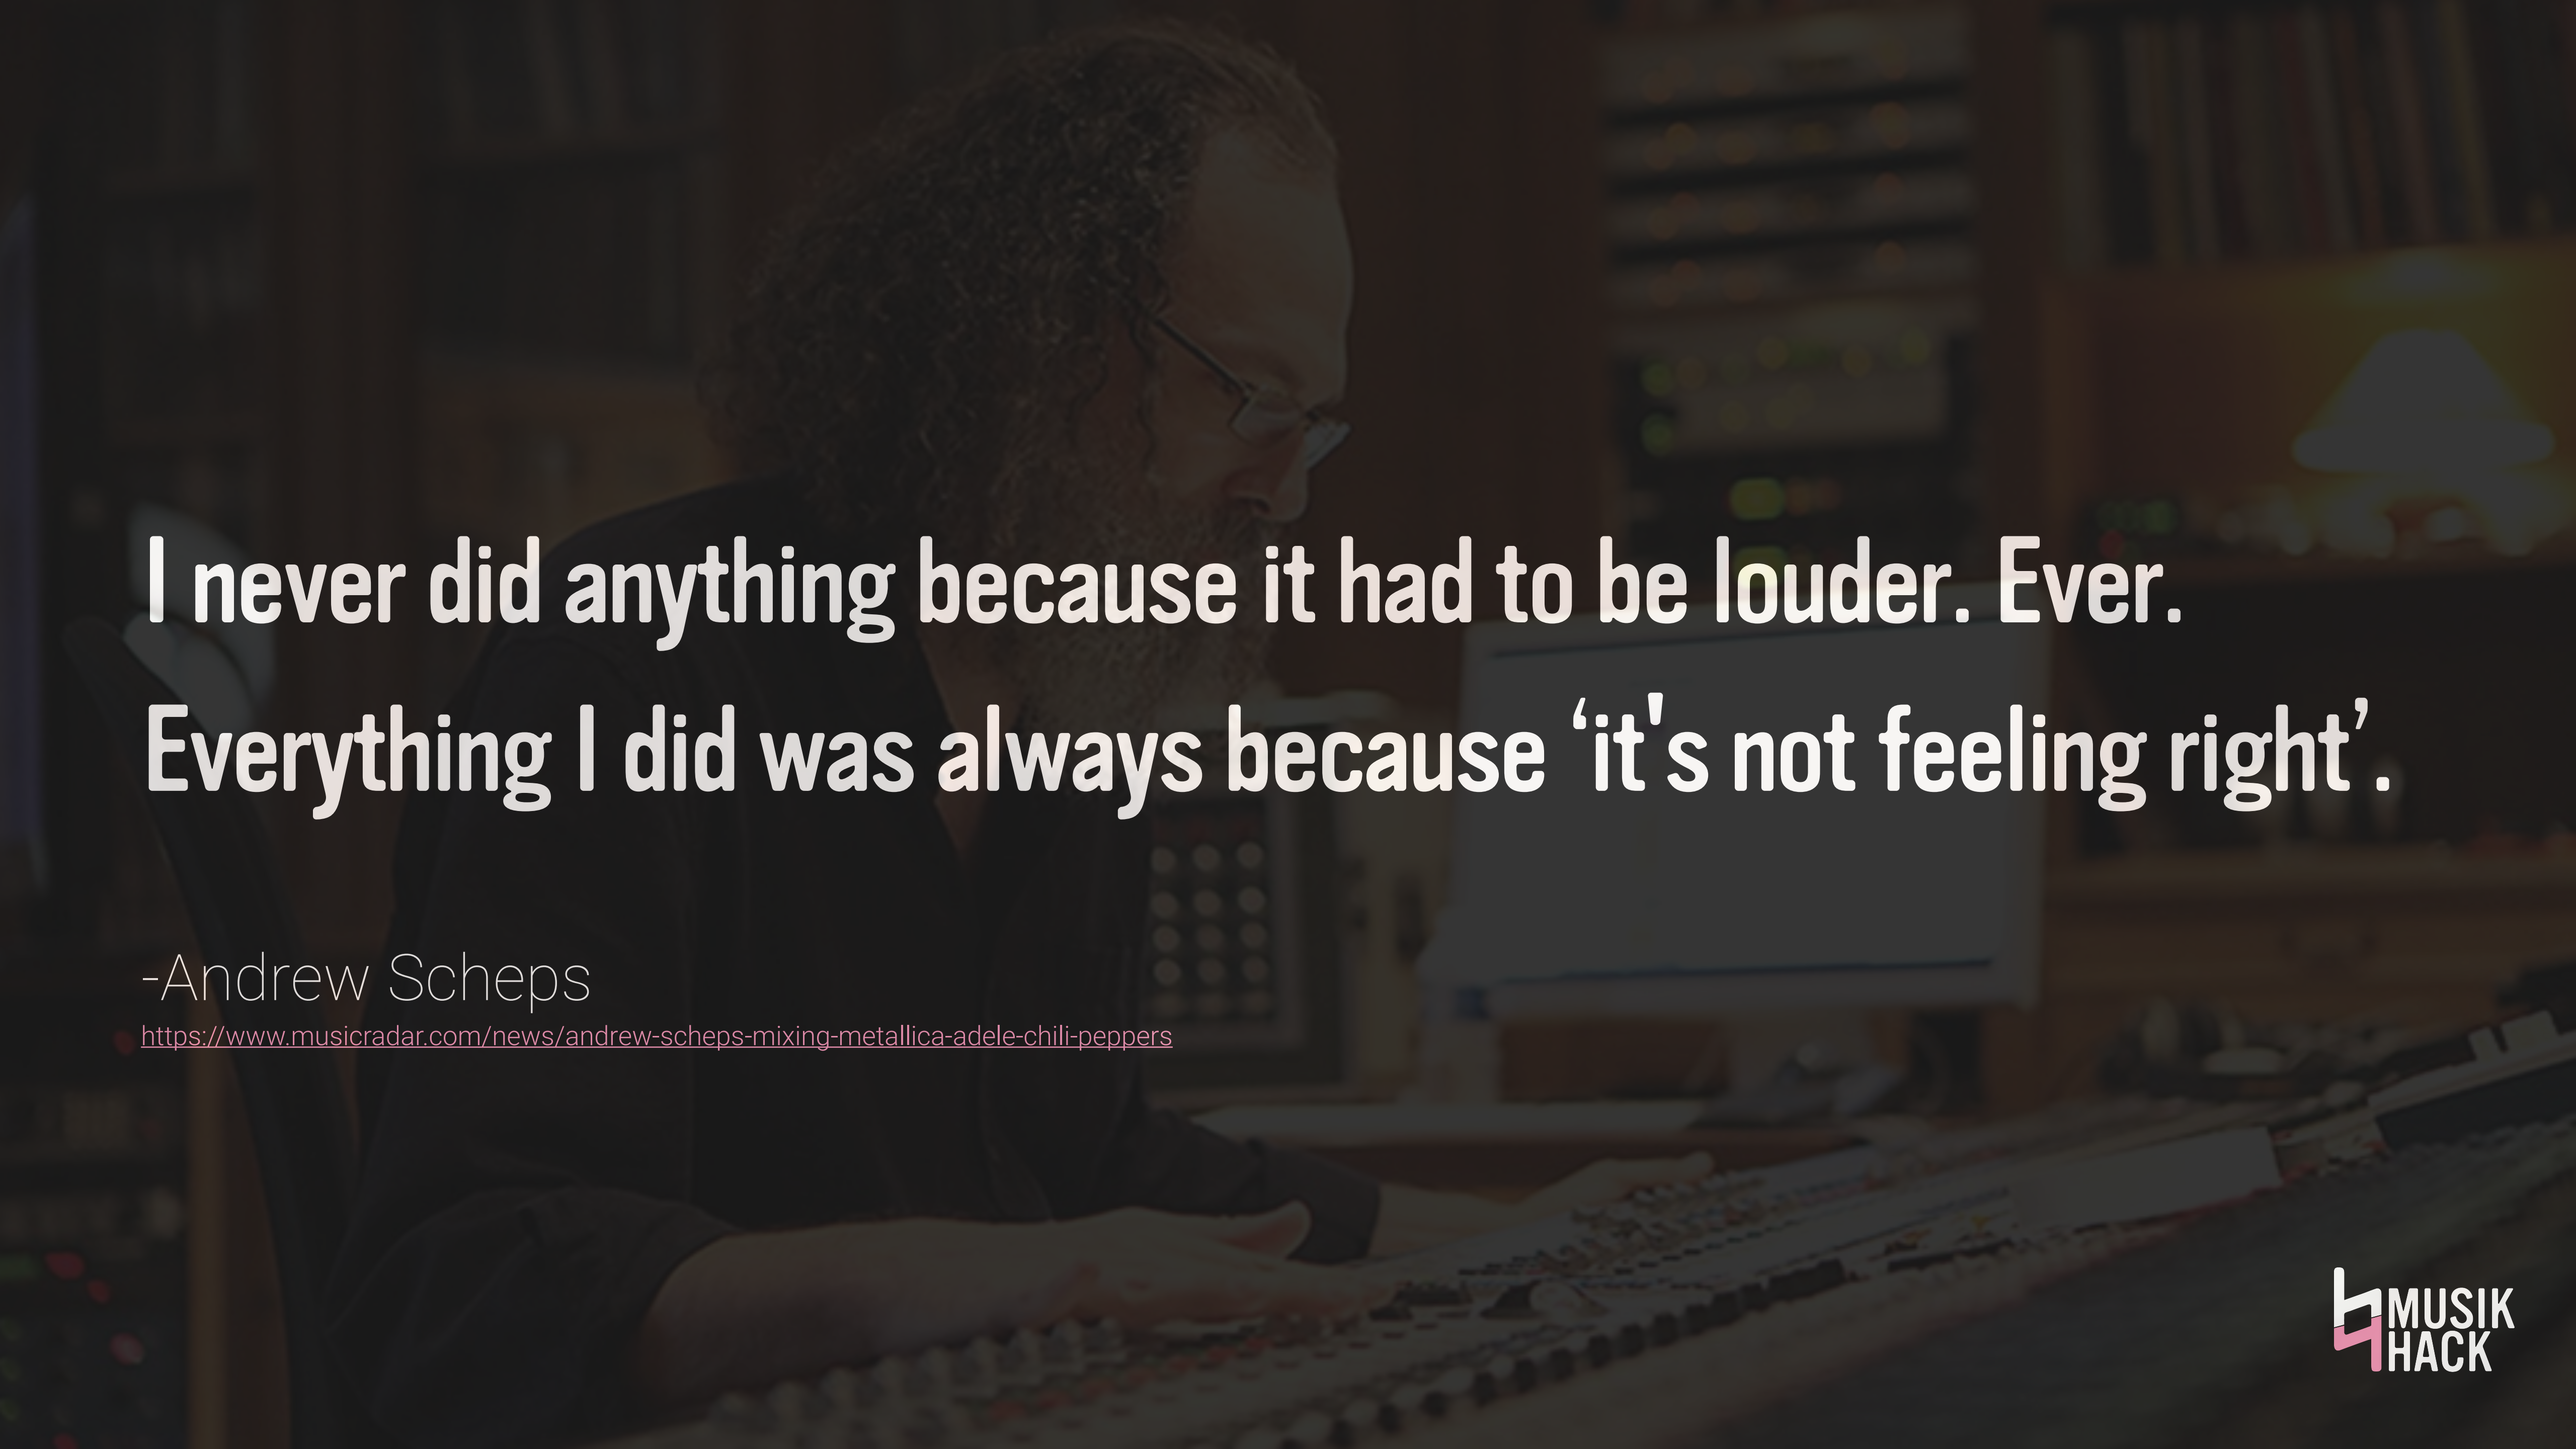

I never did anything because it had to be louder. Ever.
Everything I did was always because ‘it's not feeling right’.

-Andrew Scheps

https://www.musicradar.com/news/andrew-scheps-mixing-metallica-adele-chili-peppers

MUSIK HACK

The Modern Era: Broadcast Standards

EBU-R128, ITU-R BS.1770, AES-R7, AND SUPPLEMENTS

MUSIK HACK

Who Are These Agencies?

EBU - European Broadcast Union (Europe)

AES - Audio Engineering Society (Private / US)

ITU - International Telecommunications Union (United Nations)

MUSIK HACK

The Problem

Inconsistent loudness in program material annoys listeners

Applies to broadcasters of all types:

- TV
- Radio
- Streaming

MUSIK HACK

The Problem as Stated in EBU R 128

a) that peak normalization of audio signals has led to considerable loudness differences between programs…

b) that the resulting loudness inconsistencies… are the cause of the most viewer/listener complaints;

MUSIK HACK

The Solution per EBU 128 R

Normalize broadcasted audio to -23 LUFS (+/- 1)

Limit broadcasted audio to -1dB True Peak (+/- 0.3)

Push quiet material up (with exceptions)

Push loud material down

MUSIK HACK

Why -23?

One number to rule them all (simplicity)

Support pre-existing and varied material without normalizing up!

MUSIK HACK

Why is Normalizing up Bad?

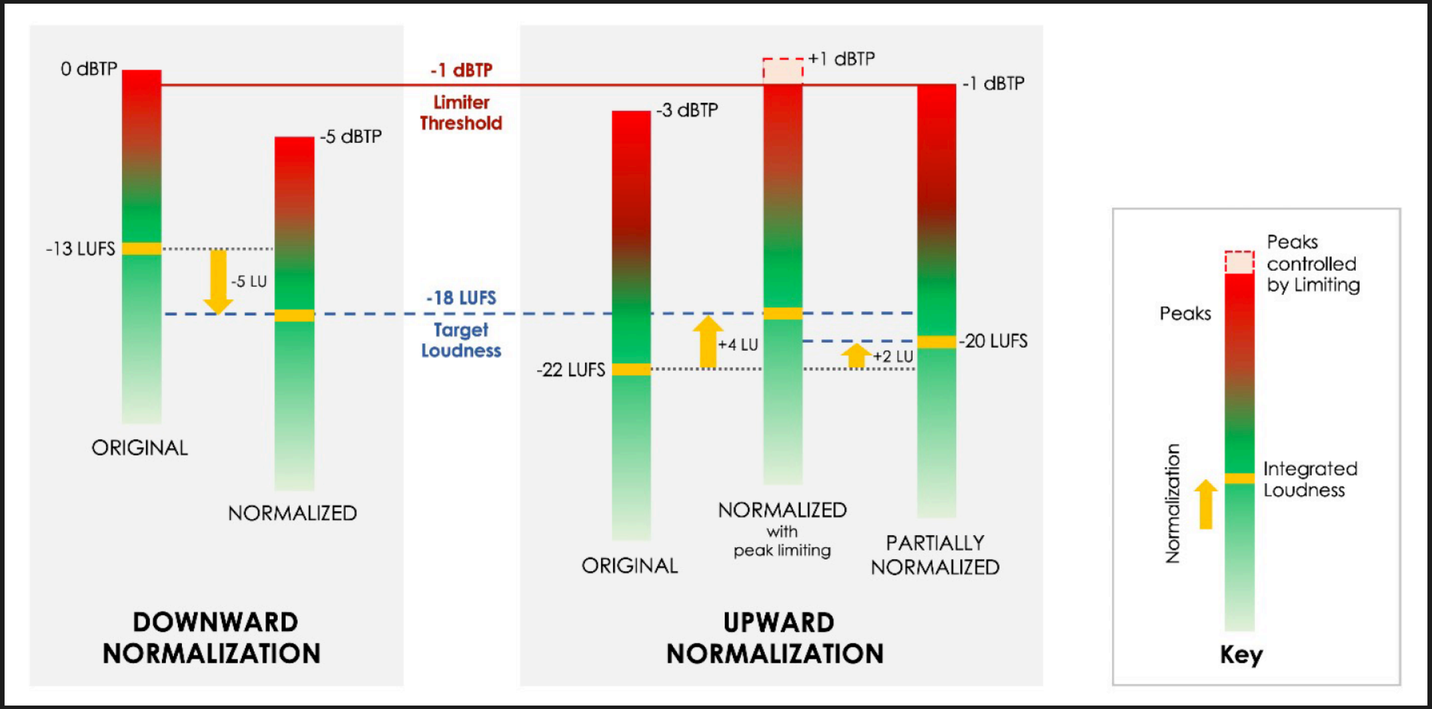

AES77-2023, *Processes for Downward and Upward Loudness Normalization*

I Still Don’t See -14 LUFS

EBU R 128 s2: Streaming Boogaloo

Personal music players: low gain, low SNR

Listening environments: high background noise

Allow -20 to -16 LUFS, TP limiter for upward normalization

Aim for -23 LUFS when players have better amplifiers

MUSIK HACK

I STILL DON'T SEE -14 LUFS!

AESTD1008 & AES77

Track vs. Album Normalization

ARE YOU HAPPY NOW?

MUSIK HACK

Adaptation

	Service	Target	Upwards?	Documented
	**Spotify**	-14	TP -1 dB no limiter	Yes
	**Spotify Loud**	-11	Yes, limiter	Yes
	**Apple Music**	-16	TP -1 dB no limiter?	No
	**YouTube**	-14	???	No

MUSIK HACK

The Hunt for Artistic Relevance

“Dynamic compression is again an artistic tool and not a loudness weapon — the **audio quality increases**!”

"**Loudness variation is an artistic tool**, and the concept of loudness normalisation according to R 128 actually **encourages more dynamic mixing**!”

“…the actual *technical* change of the audio signal level through **active normalisation to −23 LUFS** has direct influence on the *artistic* process — and in a positive way! The production side is thus relieved from fighting the ‘loudness war’ — an unfortunate result of the peak-normalisation paradigm.”

-EBU Tech 3343

GUIDELINES FOR PRODUCTION OF PROGRAMMES IN ACCORDANCE WITH EBU R 128

MUSIK HACK


THE AFTERMATH
Related searches
lufs graph
lufs for streaming
lufs meter
Related searches
lufs vs dbfs
lufs to db
lufs chart

Where we are Today

Billboard chart tracks (cross-genre) are -11 to -6 LUFS

They contain true peaks > 0 dB

They sound good (dare I say *great*), plenty of impact

The advice doesn't stick — artistically *or* commercially

One research example: https://www.izotope.com/en/learn/mastering-trends.html

MUSIK HACK

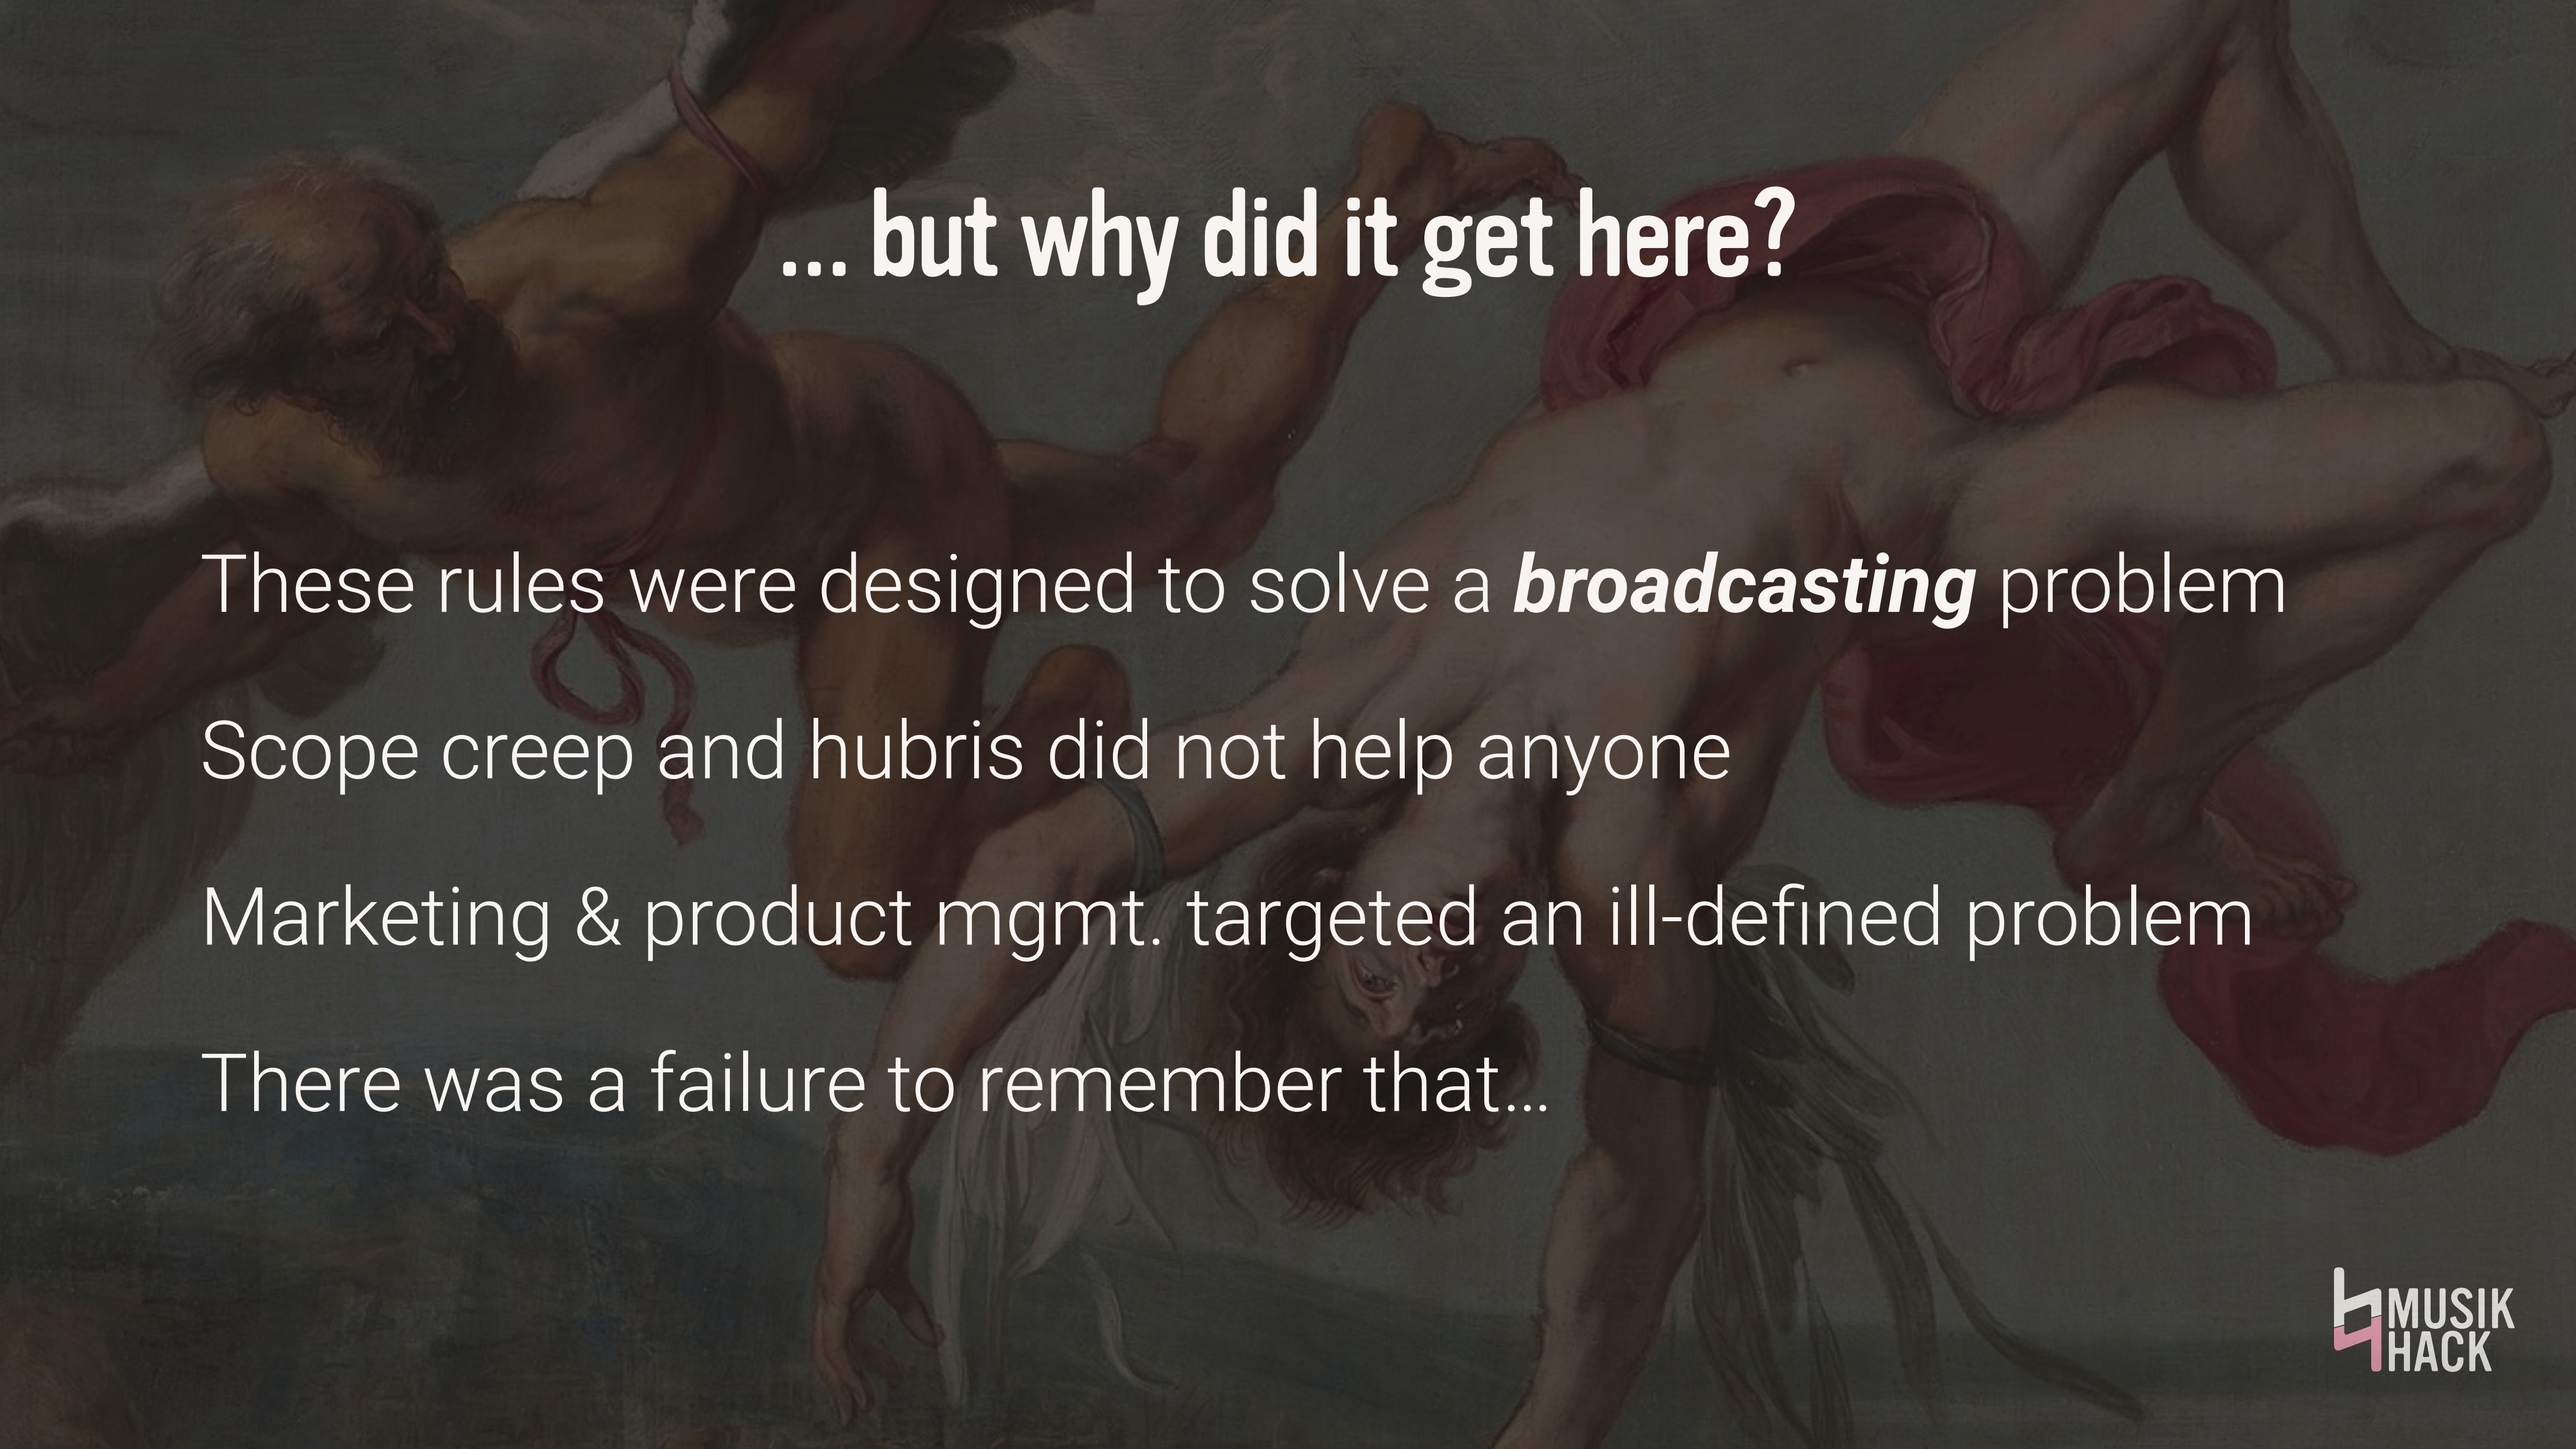

… but why did it get here?

These rules were designed to solve a ***broadcasting*** problem

Scope creep and hubris did not help anyone

Marketing & product mgmt. targeted an ill-defined problem

There was a failure to remember that…

MUSIK HACK

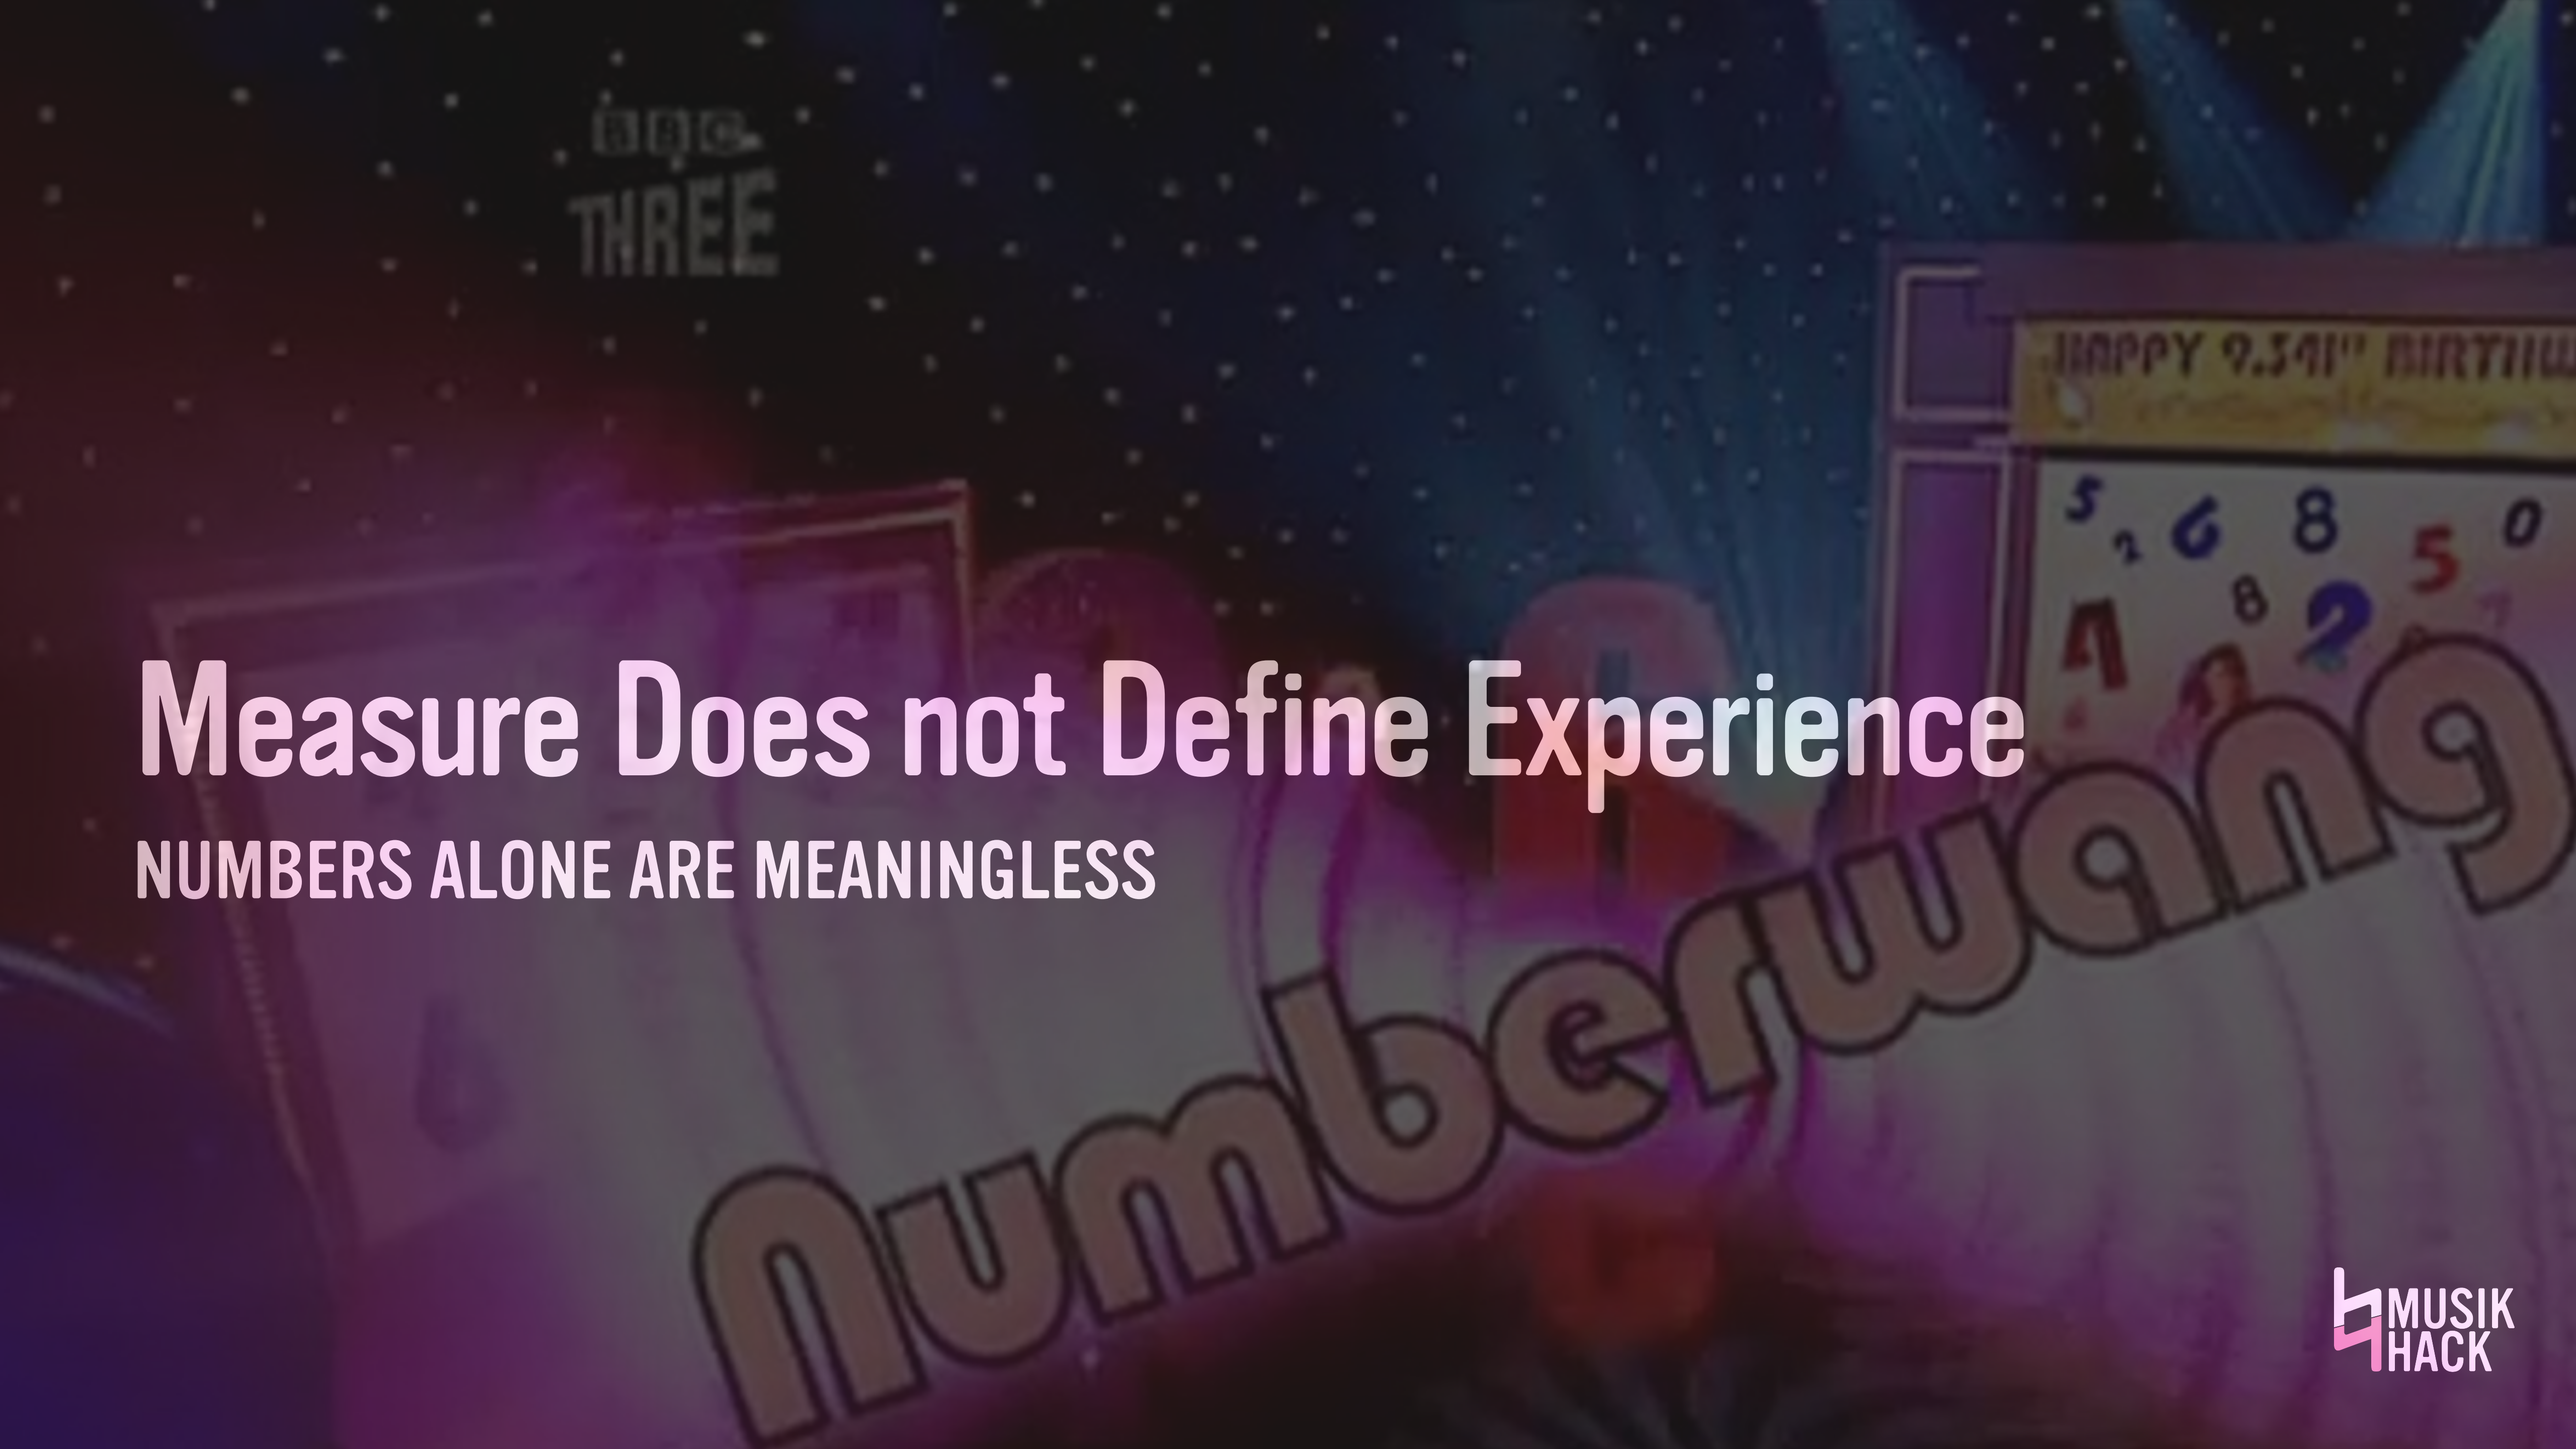

Measure Does not Define Experience

NUMBERS ALONE ARE MEANINGLESS

MUSIK HACK

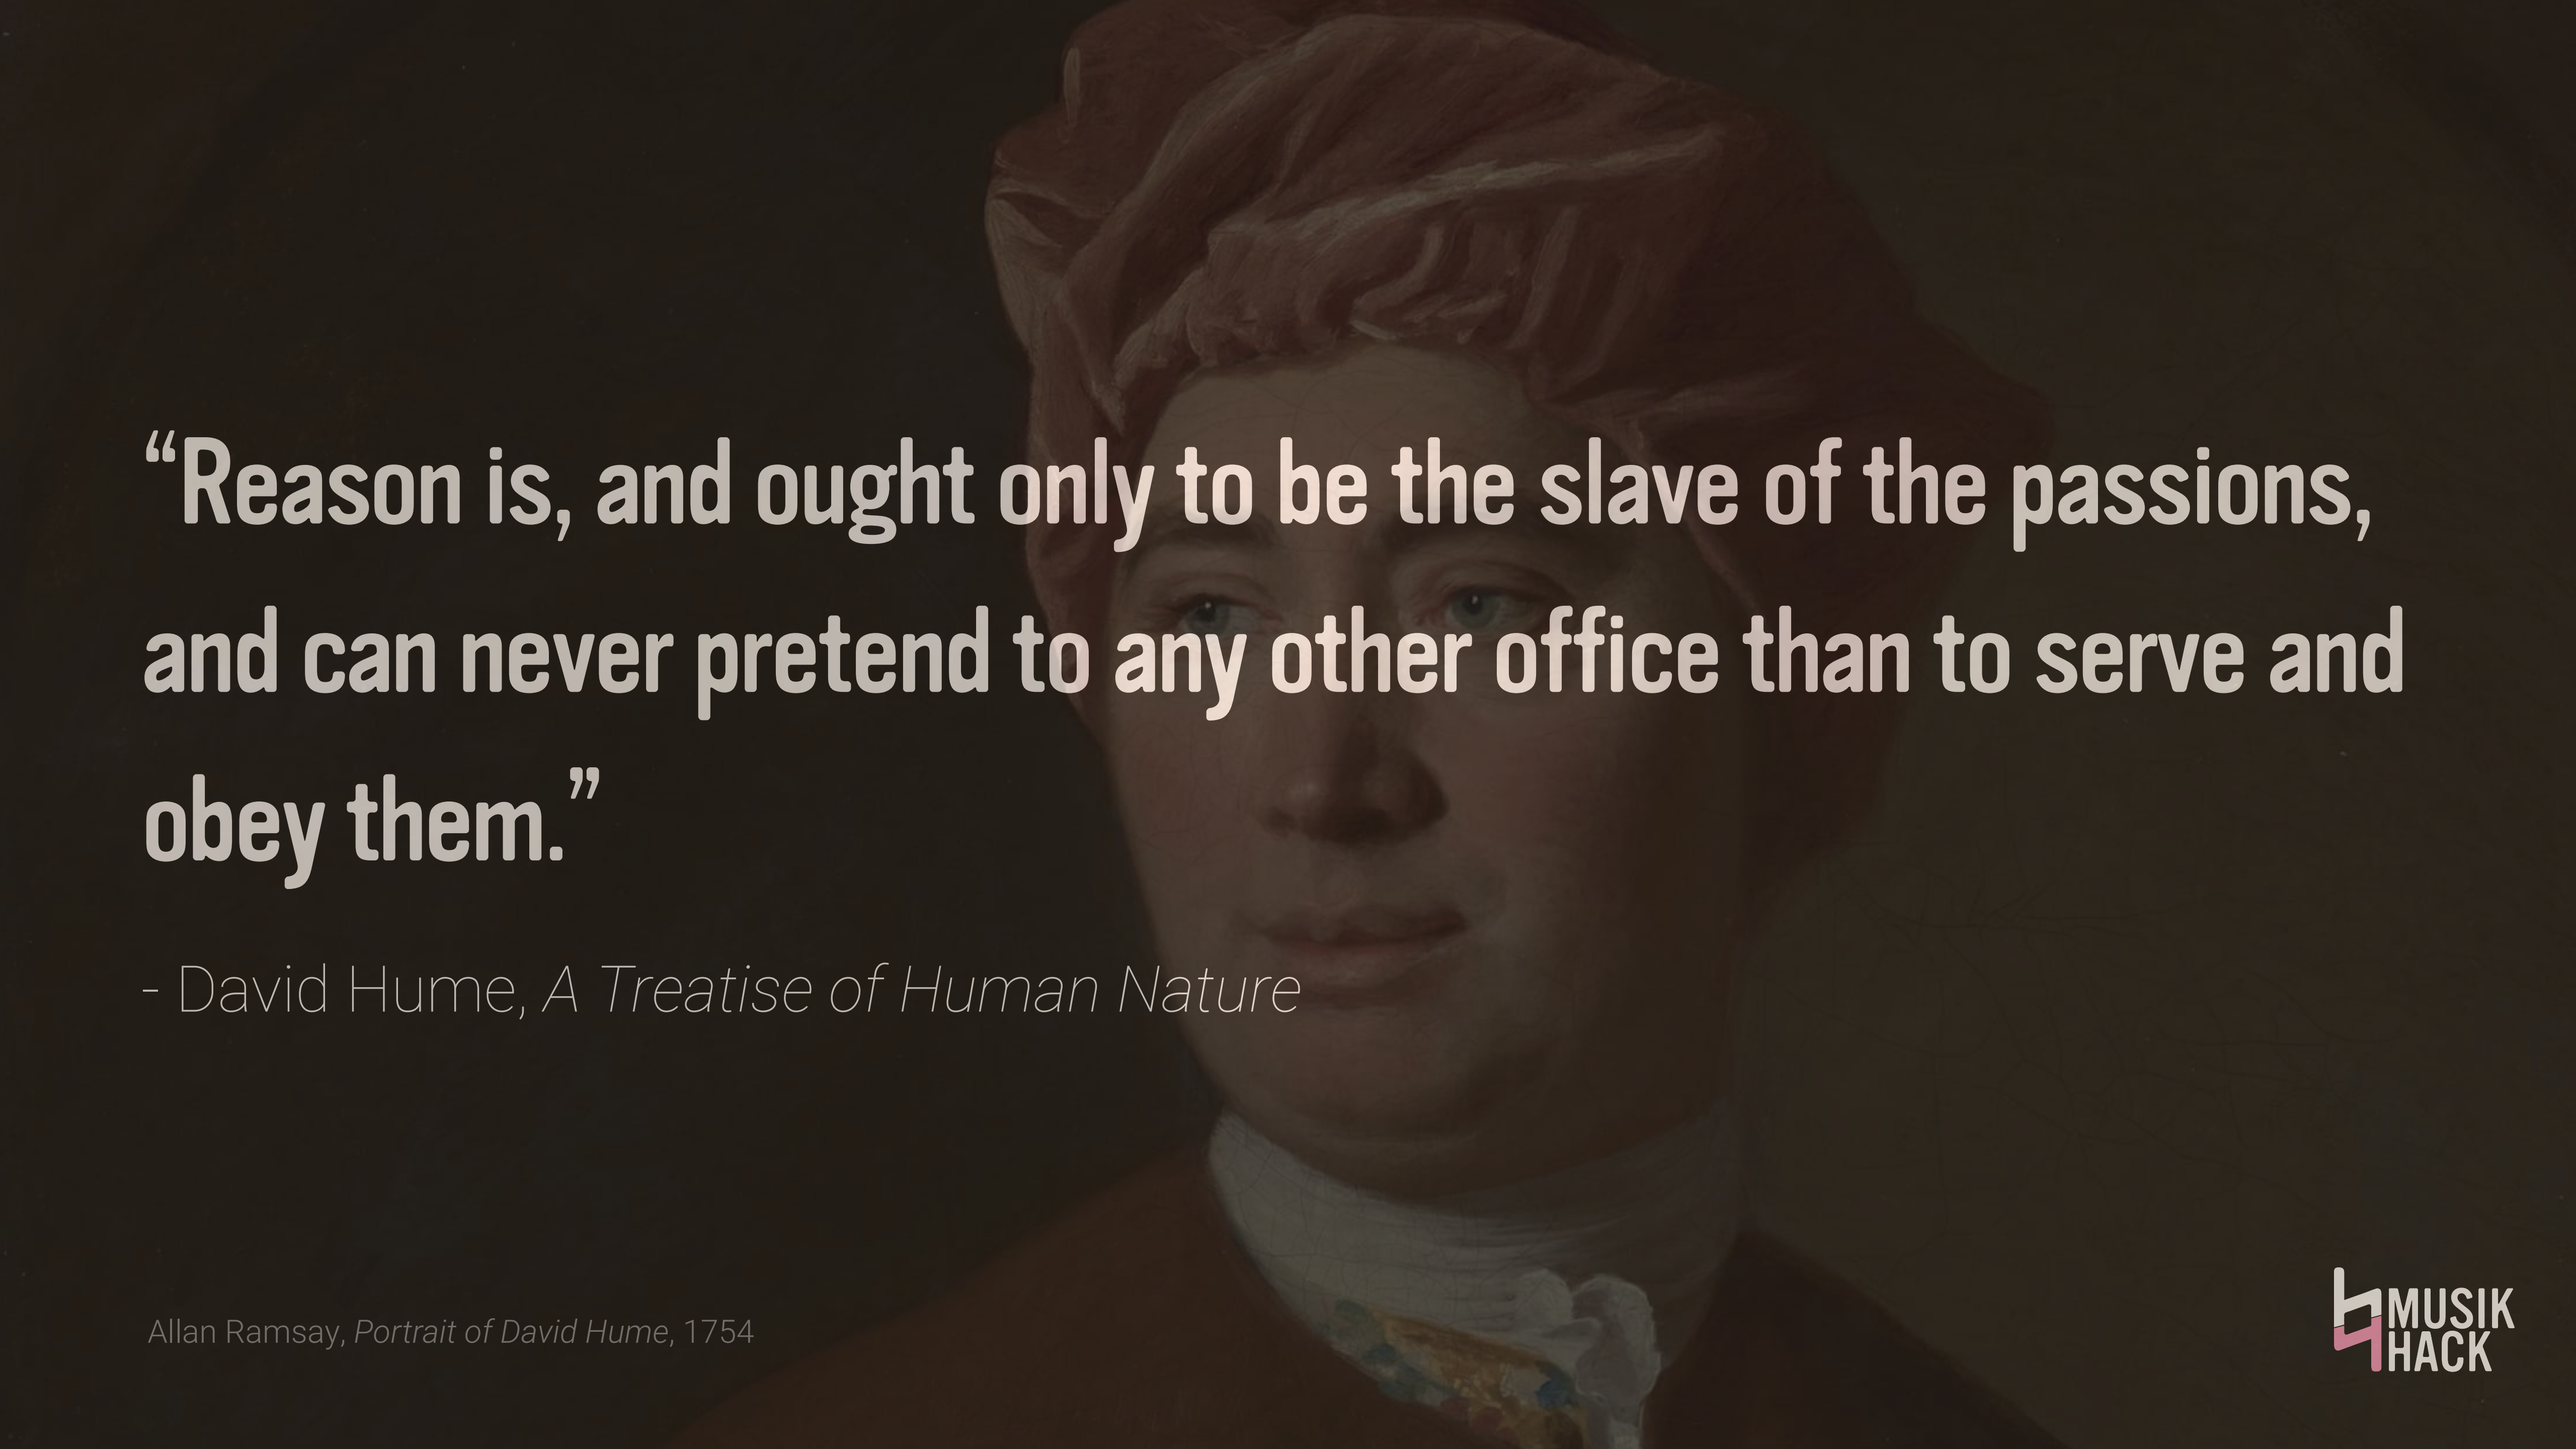

“Reason is, and ought only to be the slave of the passions, and can never pretend to any other office than to serve and obey them.”

- David Hume, *A Treatise of Human Nature*

Allan Ramsay, *Portrait of David Hume*, 1754

MUSIK HACK

Fix Complaints from the Real World

The public complained about ***inconsistent loudness!***

Did the public complain about ***poorly produced art***?

...about ***impure transients***? ...about ***limp impact***?

No, these were the passion of a (still) vocal minority!

...and they have every right to their own passions; don't impose them as gospel

MUSIK HACK

Accuracy and “Realism” != Better Production

Productions with this goal ***never*** fought the loudness war, anyway.

MUSIK HACK

True Peaks: The Untold Stories

AES R-7: True peaks only really matter for highs/transients

(Transients are short noise bursts, so already mask peaks heavily)

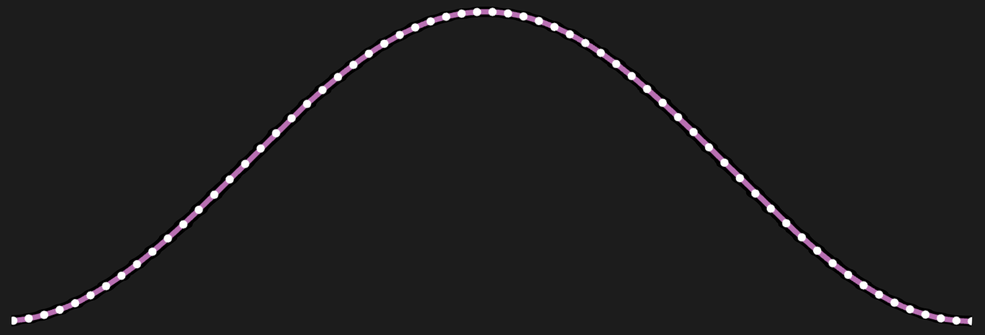

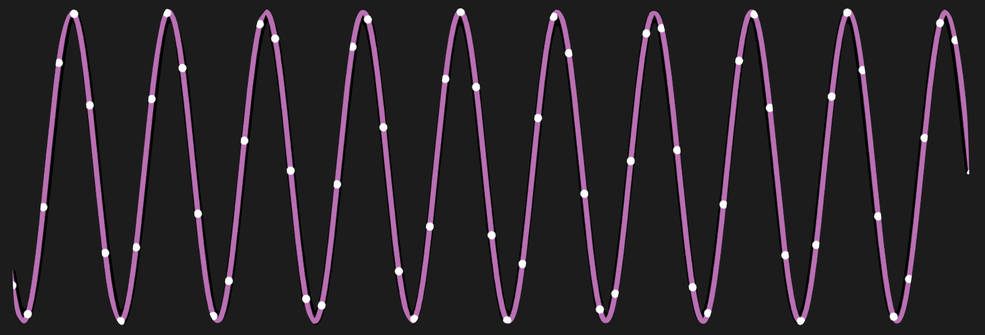

MUSIK HACK

True Peaks: The Untold Stories contd.

AES R-7:

TP meters *still* have a max under-read up to 0.6 dB

AES 77-2023-F:

"...true peak limiting more than about 1dB may ***produce more audible artifacts than simply letting audio material clip***"

MUSIK HACK

True Peaks: Codecs

Lossy codecs don’t clip on encode

Lossy codecs don’t inherently clip on decode

Volume normalization with FP decoder = HEADROOM

Are occasional clips more audible than the “lossiness”?

You can always listen…

https://www.apple.com/apple-music/apple-digital-masters/

MUSIK HACK

In Conclusion, Listen

THE BEST ADVICE IS TRITE

MUSIK HACK

Dynamic Range, PLR

Not all dynamic range compression is the same

Dynamic range by itself is not a virtue

Perceived dynamics/impact > measured dynamics

MUSIK HACK

Loudness is Robust & Predictable

Sounds great; does not necessarily lack impact

Better listenability in noisy environments

No upwards normalization/standards safe

Less extreme peaks, lower amplification required

Lower amplification == less analog nonlinearity

More predictable & listenable in real world systems

MUSIK HACK

No man is wise at all times, or is without his blind side.

-Desiderius Erasmus, In Praise of Folly

MUSIK HACK

THANK YOU!

MUSIK HACK