Music 209
Advanced Topics in Computer Music
Lecture 2 – Splicing

SYNFUL

ORCHESTRA

Special guest: Eric Lindemann

2006-1-26

Professor David Wessel (with John Lazzaro)

(cnmat.berkeley.edu/~wessel, www.cs.berkeley.edu/~lazzaro)

www.cs.berkeley.edu/~lazzaro/class/music209

Last time: Course Introduction ...

Why pianos were easy ... compared to a violin.

	Piano	Violin
Articulations	**One**	**Many**
Expression During Sustain	**No**	**Yes**
Legato and Portamento	**No**	**Yes**

Today: Automatic Splicing

- **Flowchart for concatenation.**
- **Choosing good matches.**
- **Doing good splices.**

Recall: Legato Concatenation

Actually involves a series of decisions ...

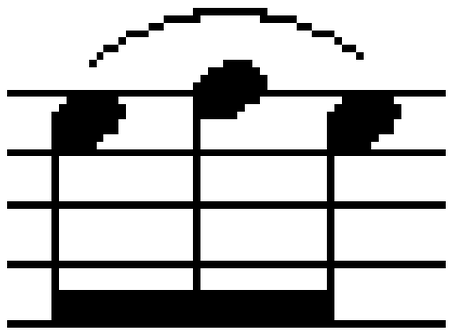

	Piano	Violin
Legato and Portamento	No	Yes

"Online" stitching of a legato run from 3 samples in a library

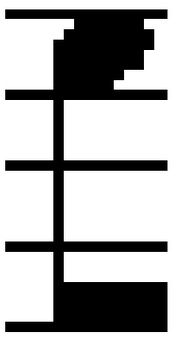

splice to

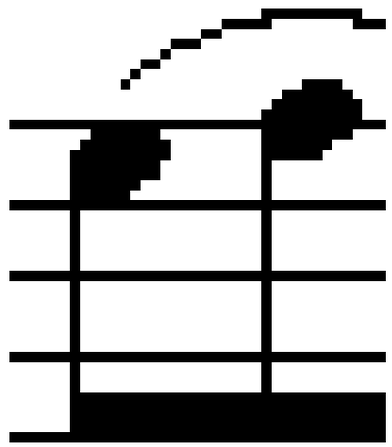

splice to

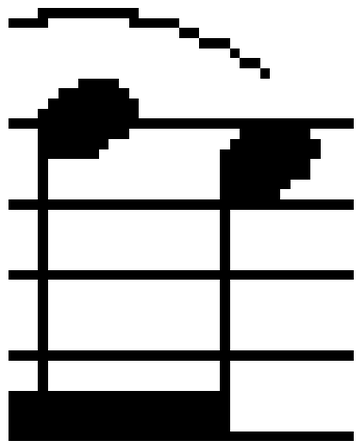

Sample #1: isolated E

Sample #2: E to F interval played legato

Sample #3: F to E interval played legato

Starting sample

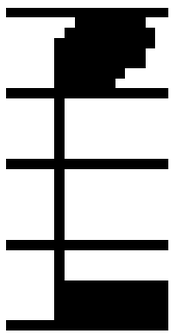

Select candidate samples from db

Candidates

#1

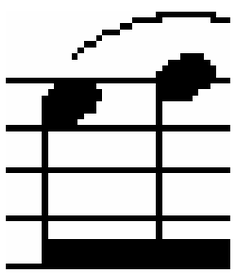

#2

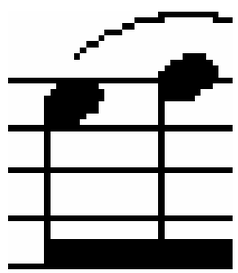

#3

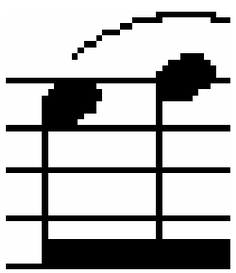

Any good matches?

N

Modify a candidate to be good enough

#3(mod)

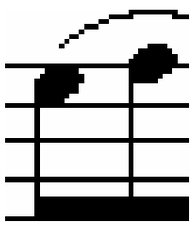

Do the splice

#3(mod)

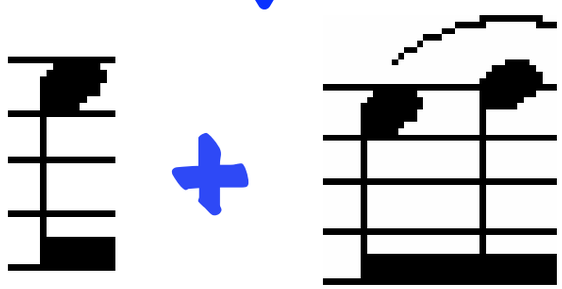

Y

Choose best candidate

#2

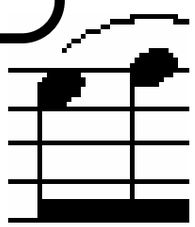

Do the splice

#2

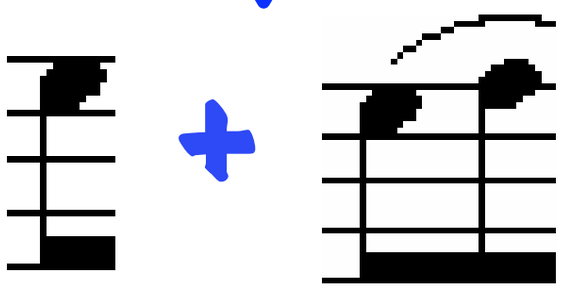

Today, we focus on match evaluation and splicing.

Gray regions are topics of later lectures.

Select candidate samples from db

Any good matches?

N

Modify a candidate to be good enough

#3(mod)

Do the splice

#3(mod)

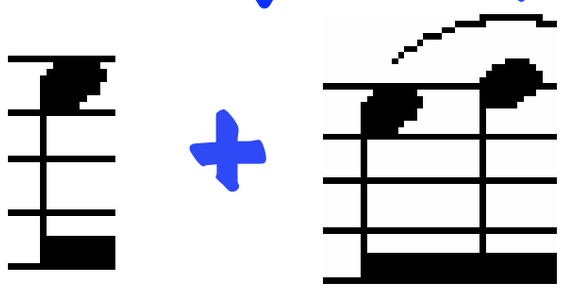

Y

Choose best match

#2

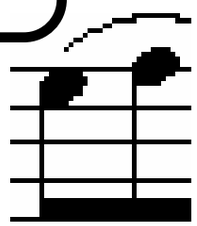

Do the splice

#2

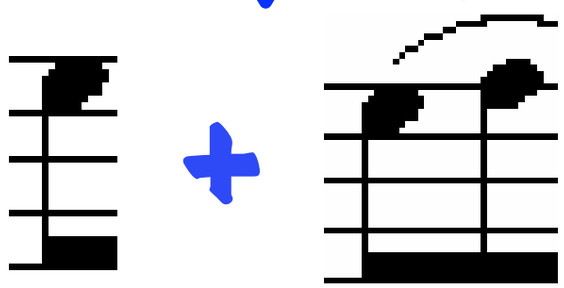

Candidates

#1

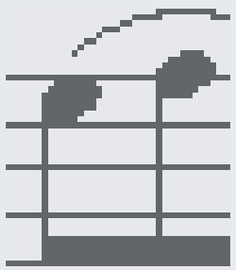

#2

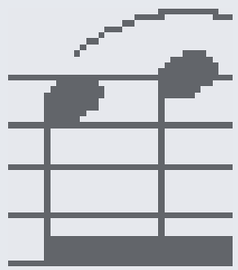

#3

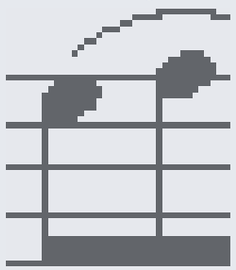

First up, match evaluation ...

Select candidate samples from db

Any good matches?

N

Modify a candidate to be good enough

#3(mod)

Do the splice

#3(mod)

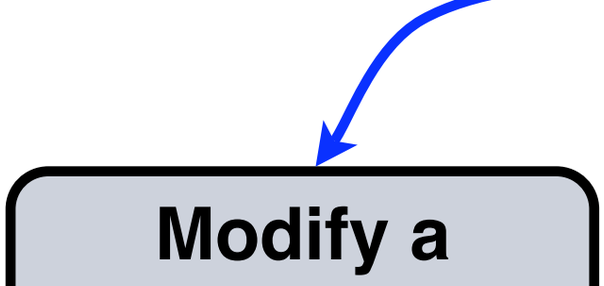

Y

Choose best match

#2

Do the splice

#2

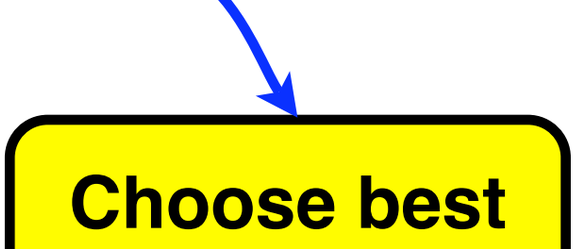

Candidates

#1

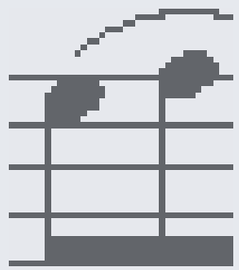

#2

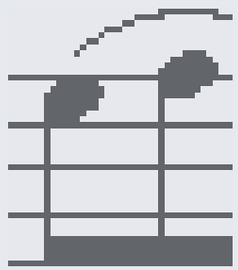

#3

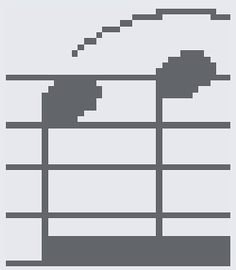

What makes a good (or the best) match?

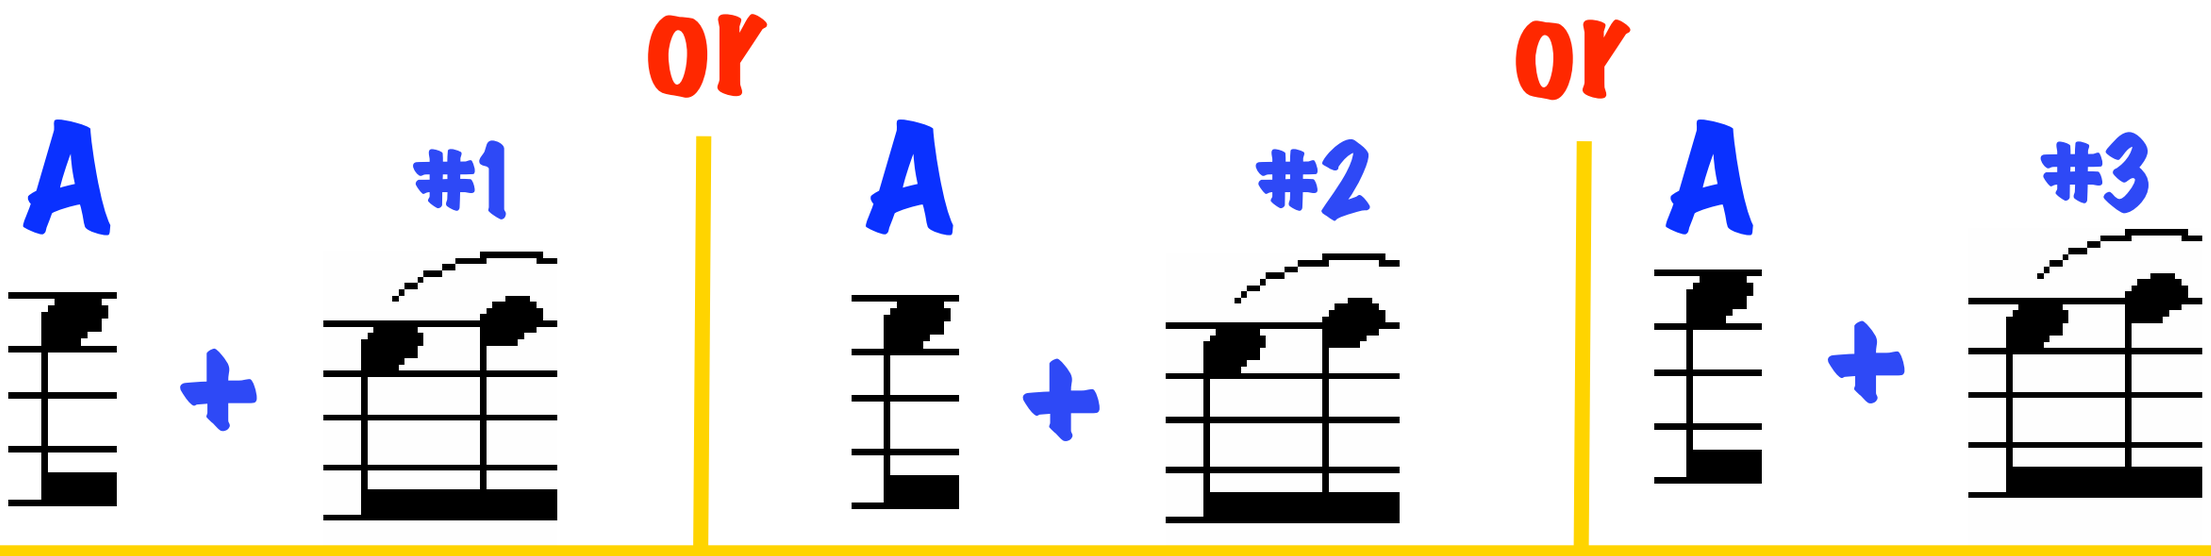

Given samples A and B, we define a metric f(A, B) of concatenation quality. Compare f(A, #1), f(A, #2), f(A, #3) to find the best.

Compare best f() against an absolute standard to test for good enough.

Our metric depends on the application.

- **Transparency metric.** The end of A and the beginning of B are selected to be **nearly identical.** We are looking for a splice that sounds transparent (i.e. not noticable).

- **Fusion metric.** The end of A and the start of B have **different timbres** (example: A is the "ta" onset of a trumpet, B is a sustained sound). We are looking for **perceptual fusion** across the splice.

Other applications may need different metrics.

- **Rhythm metric.** A and B are drum patterns. Our splice should be **rhythmically smooth**.

Early Transparent Splicing: Sustain loops

1984,
Emulator II,
$8000.
Not the first
sampler, but ...

Sufficiently inexpensive to bring sampling to the working pro.

512 KB of RAM -- 17s @ 27 kHz (8-bit companded)

To use many samples across the keyboard, could not afford the "20-40s per piano key" approach!

Solution: "loop" a 1-100 ms sustained portion of the piano waveform to play over and over until key release.

Transparent (self) concatenation.

How it works:

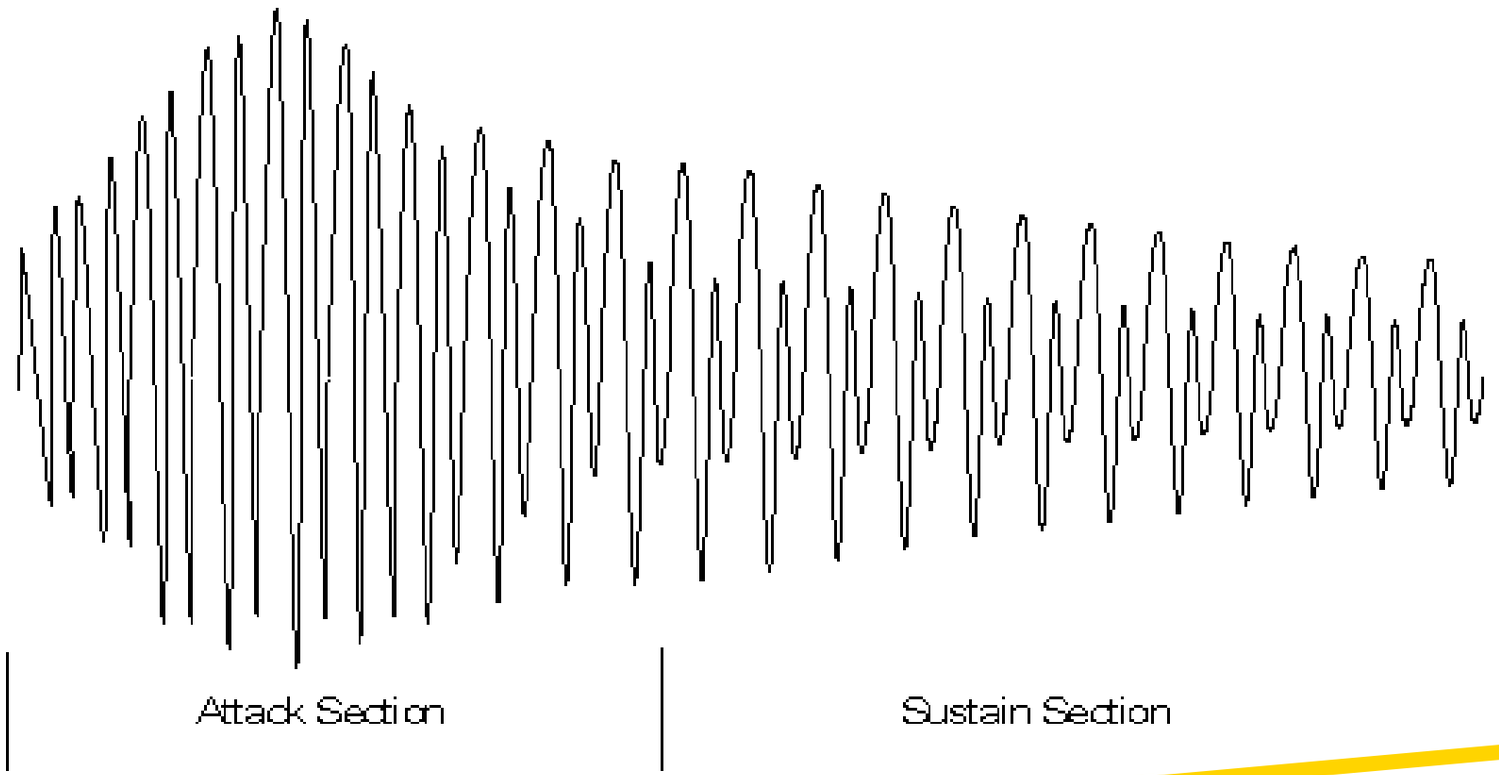

Original instrument recording.

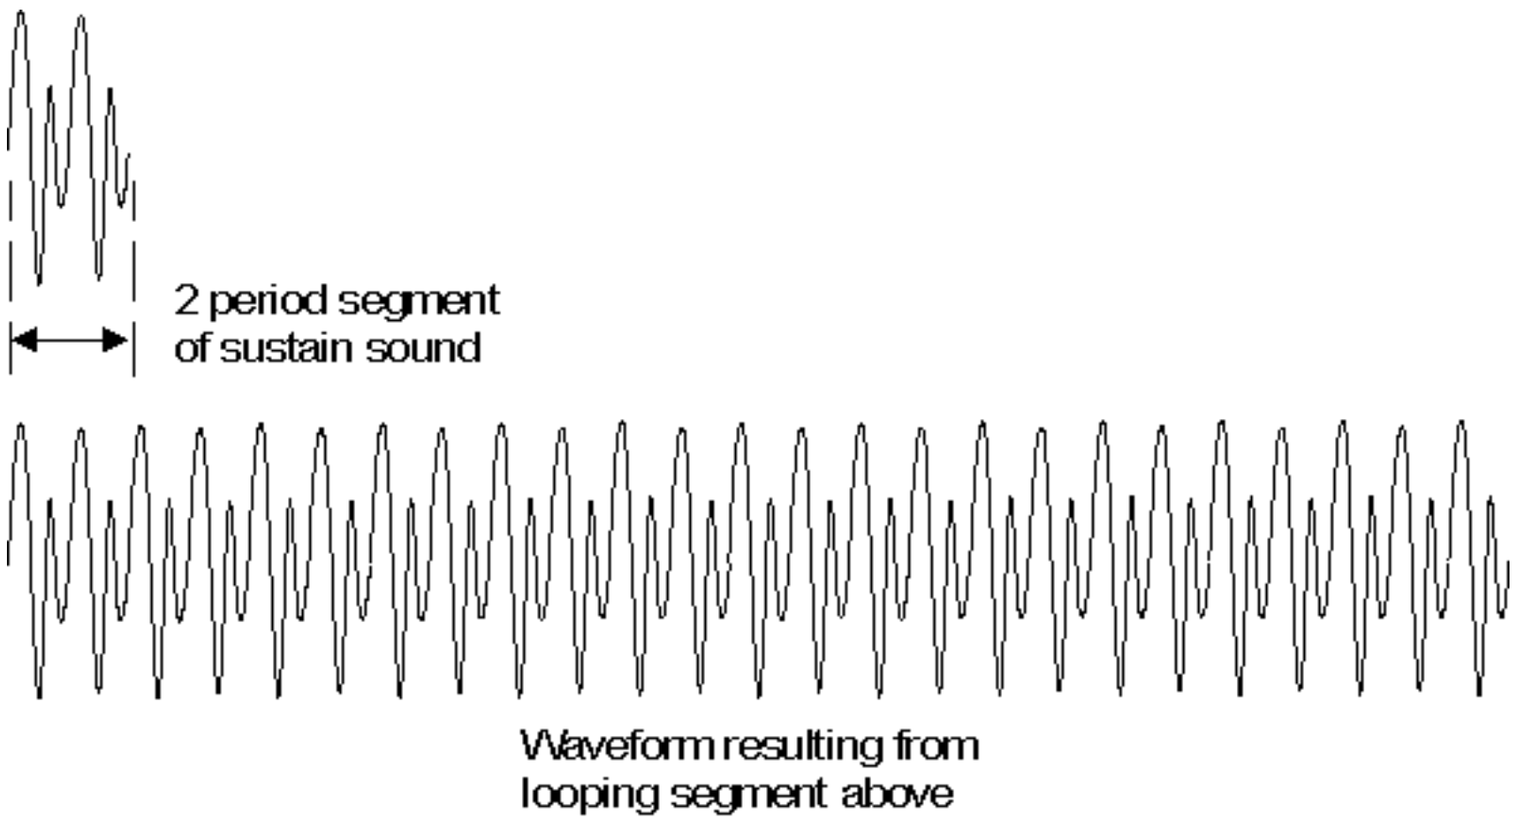

Isolating a part of the sustained section that will loop transparently.

An art and a science ...

Images from Jim Heckrock tutorial on Harmony Central

Concatenate attack and loops, envelope

Original attack ...

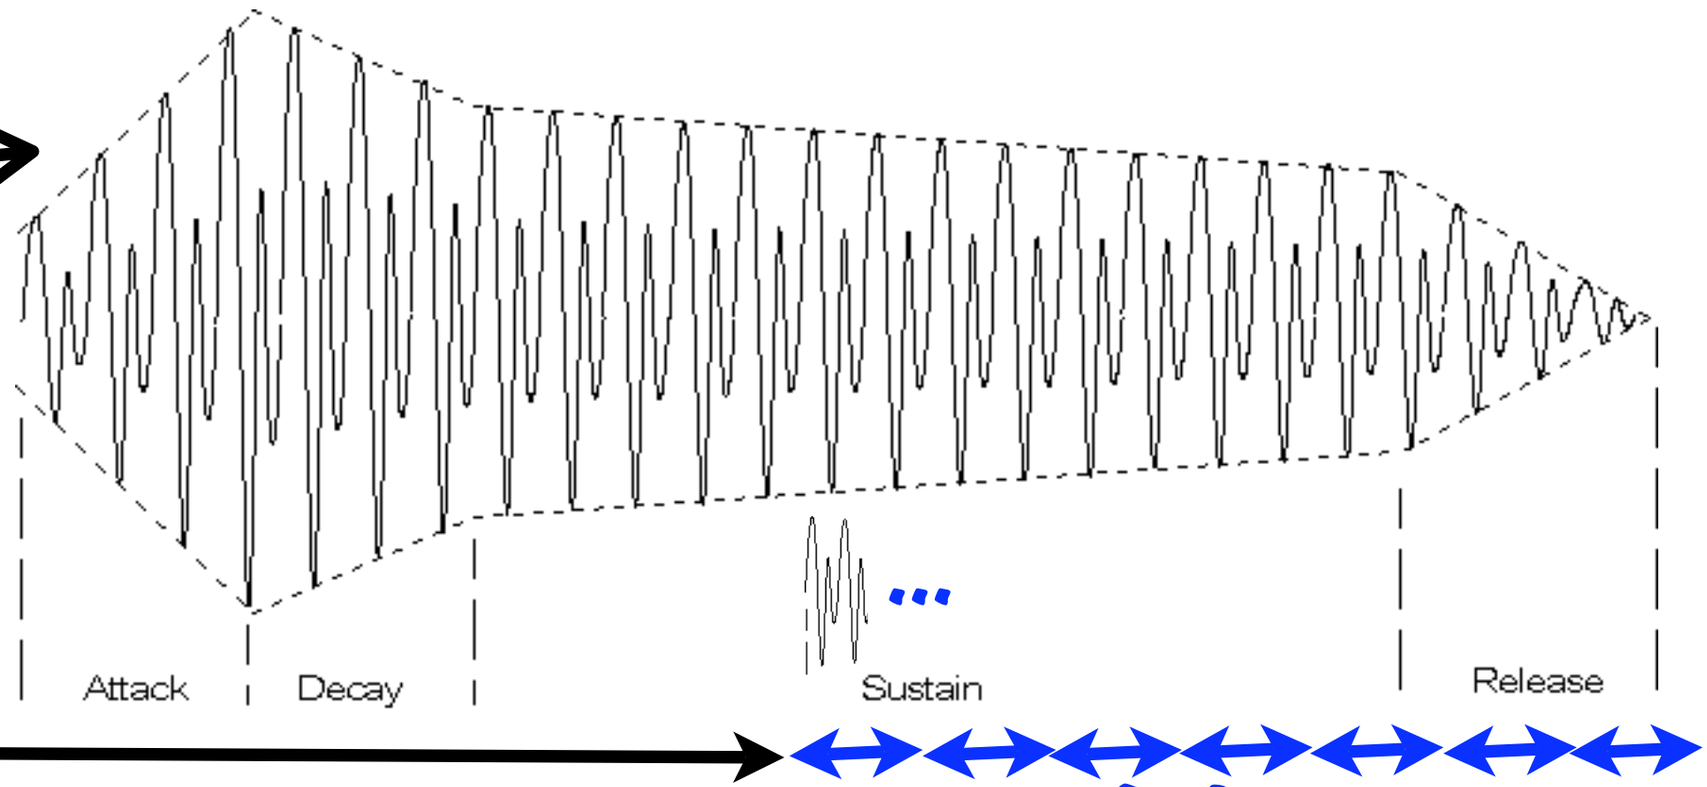

Splices

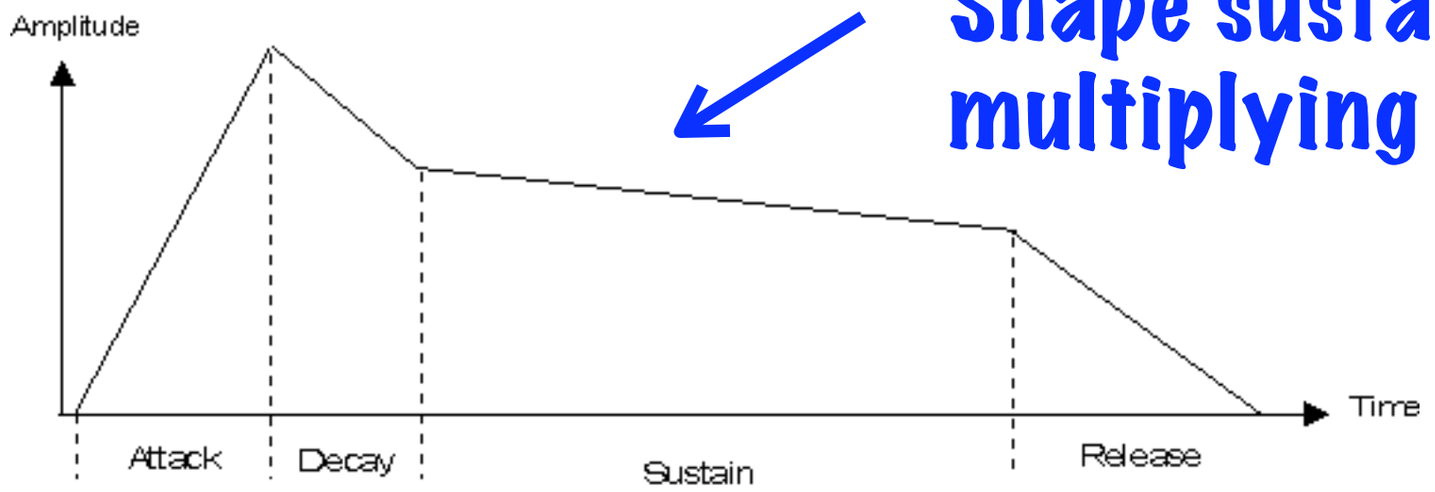

Shape sustain and release by multiplying by this contour.

Images from Jim Heckrock tutorial on Harmony Central

Why our problem is harder ...

Instead of doing many splices of a short segment of one recording ...

2 period segment of sustain sound

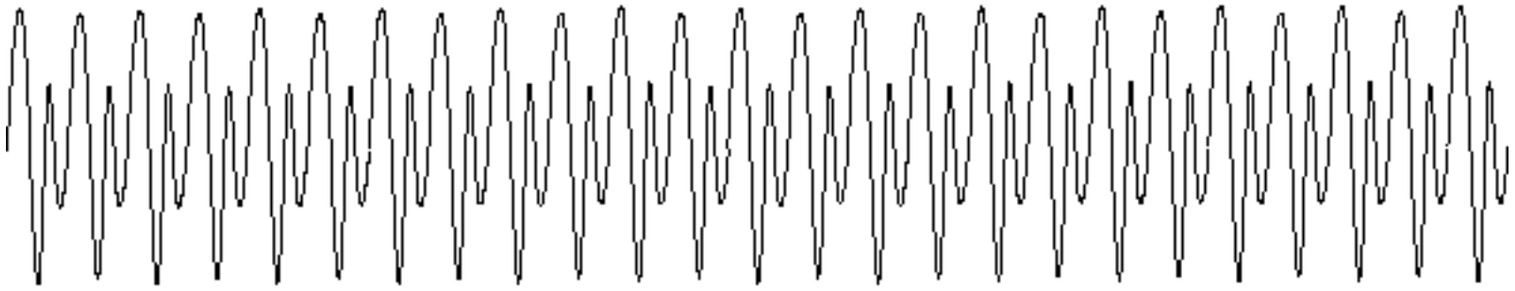

Waveform resulting from looping segment above

we are doing one splice between segments of two recordings ...

Closer to us: Looping animated sounds

Ensoniq VFX ROM string ensemble sample.

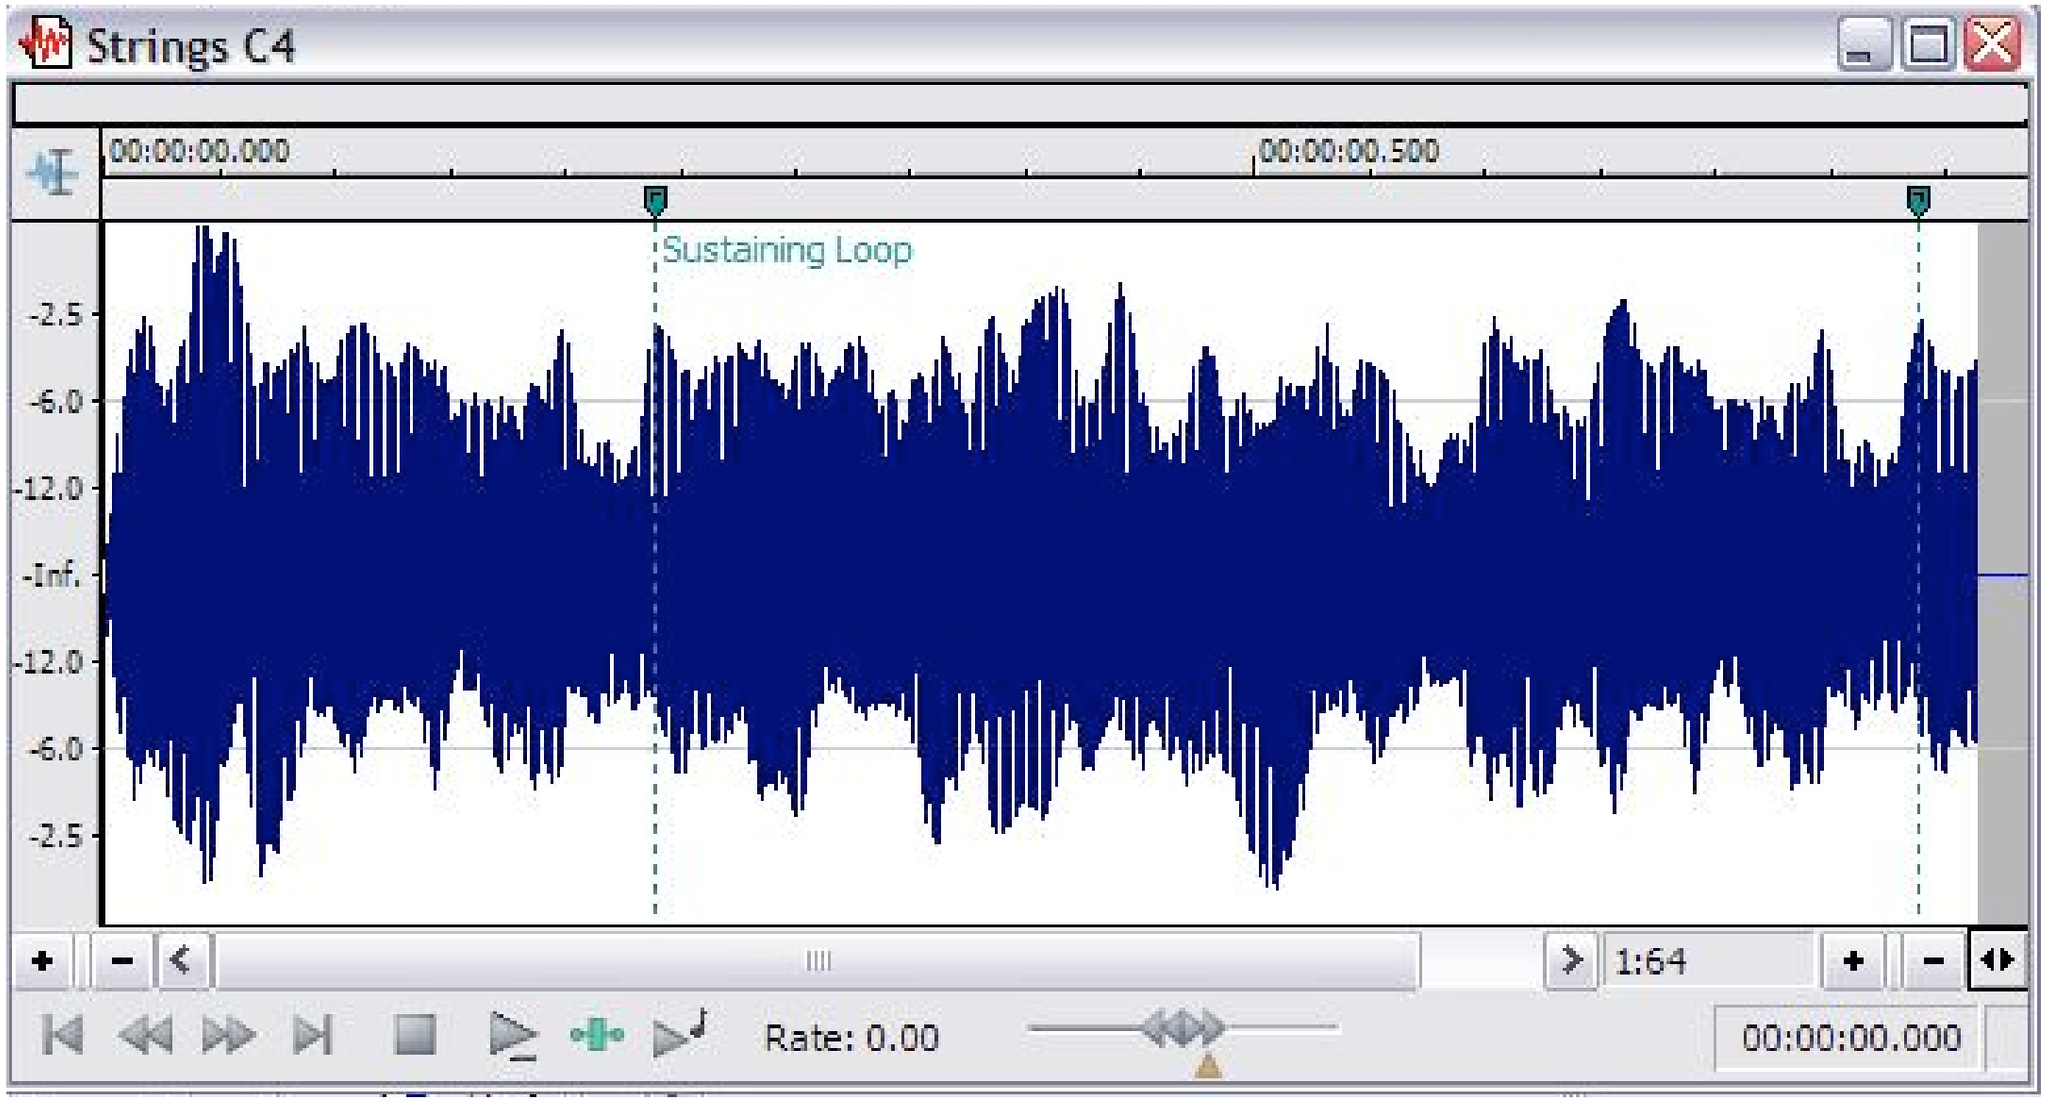

500 ms loop of an "animated" sound whose timbre is constantly changing ... looping a sound like this (manually) is actually possible!

Image from Tweakheadz.com sampling tutorial.

Recall: Concatenative transparency metric

Transparency metric. The end of A and the beginning of B are selected to be **nearly identical.** We are looking for a splice that sounds transparent (i.e. not noticable).

What makes a transparent splice?

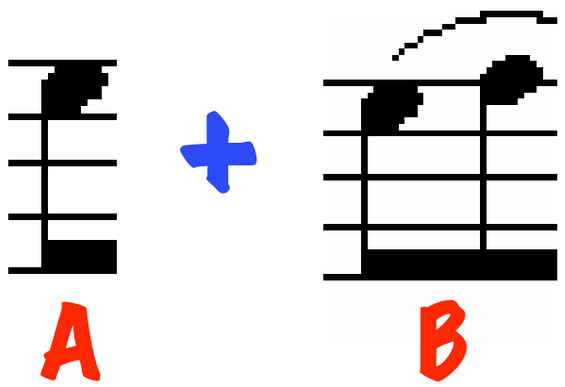

No waveform discontinuity at the splice point. **Easy to handle in the "do the splice" algorithm.**

Harder: The end of A and the start of B should have ...

- **Similar loudness.**
  - Absolute & delta: amplitude envelope, tremolo.
- **Similar spectral shape.**
  - Absolute & delta: spectral motion across splice.
- **Similar pitch.**
  - Absolute & delta: vibrato and pitch bends.

Next topic -- if we have a good enough match, how to we actually do the splice?

Select candidate samples from db

Candidates

#1

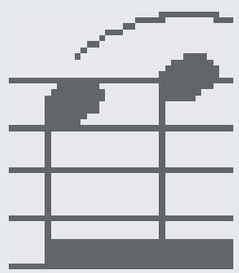

#2

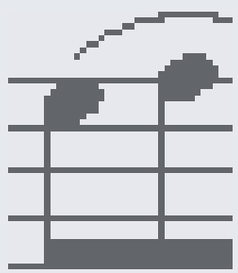

#3

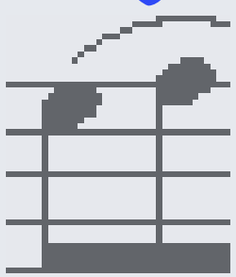

Any good matches?

N

Modify a candidate to be good enough

#3(mod)

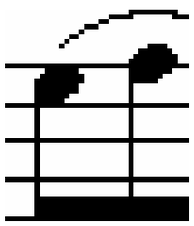

Do the splice

#3(mod)

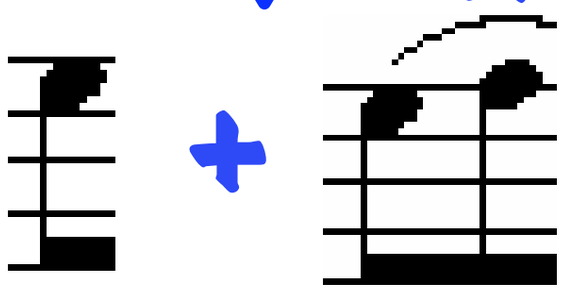

Y

Choose best match

#2

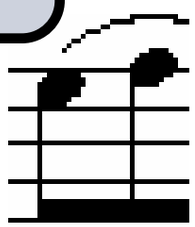

Do the splice

#2

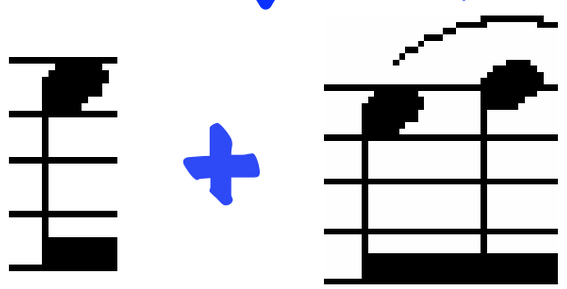

Crossfades

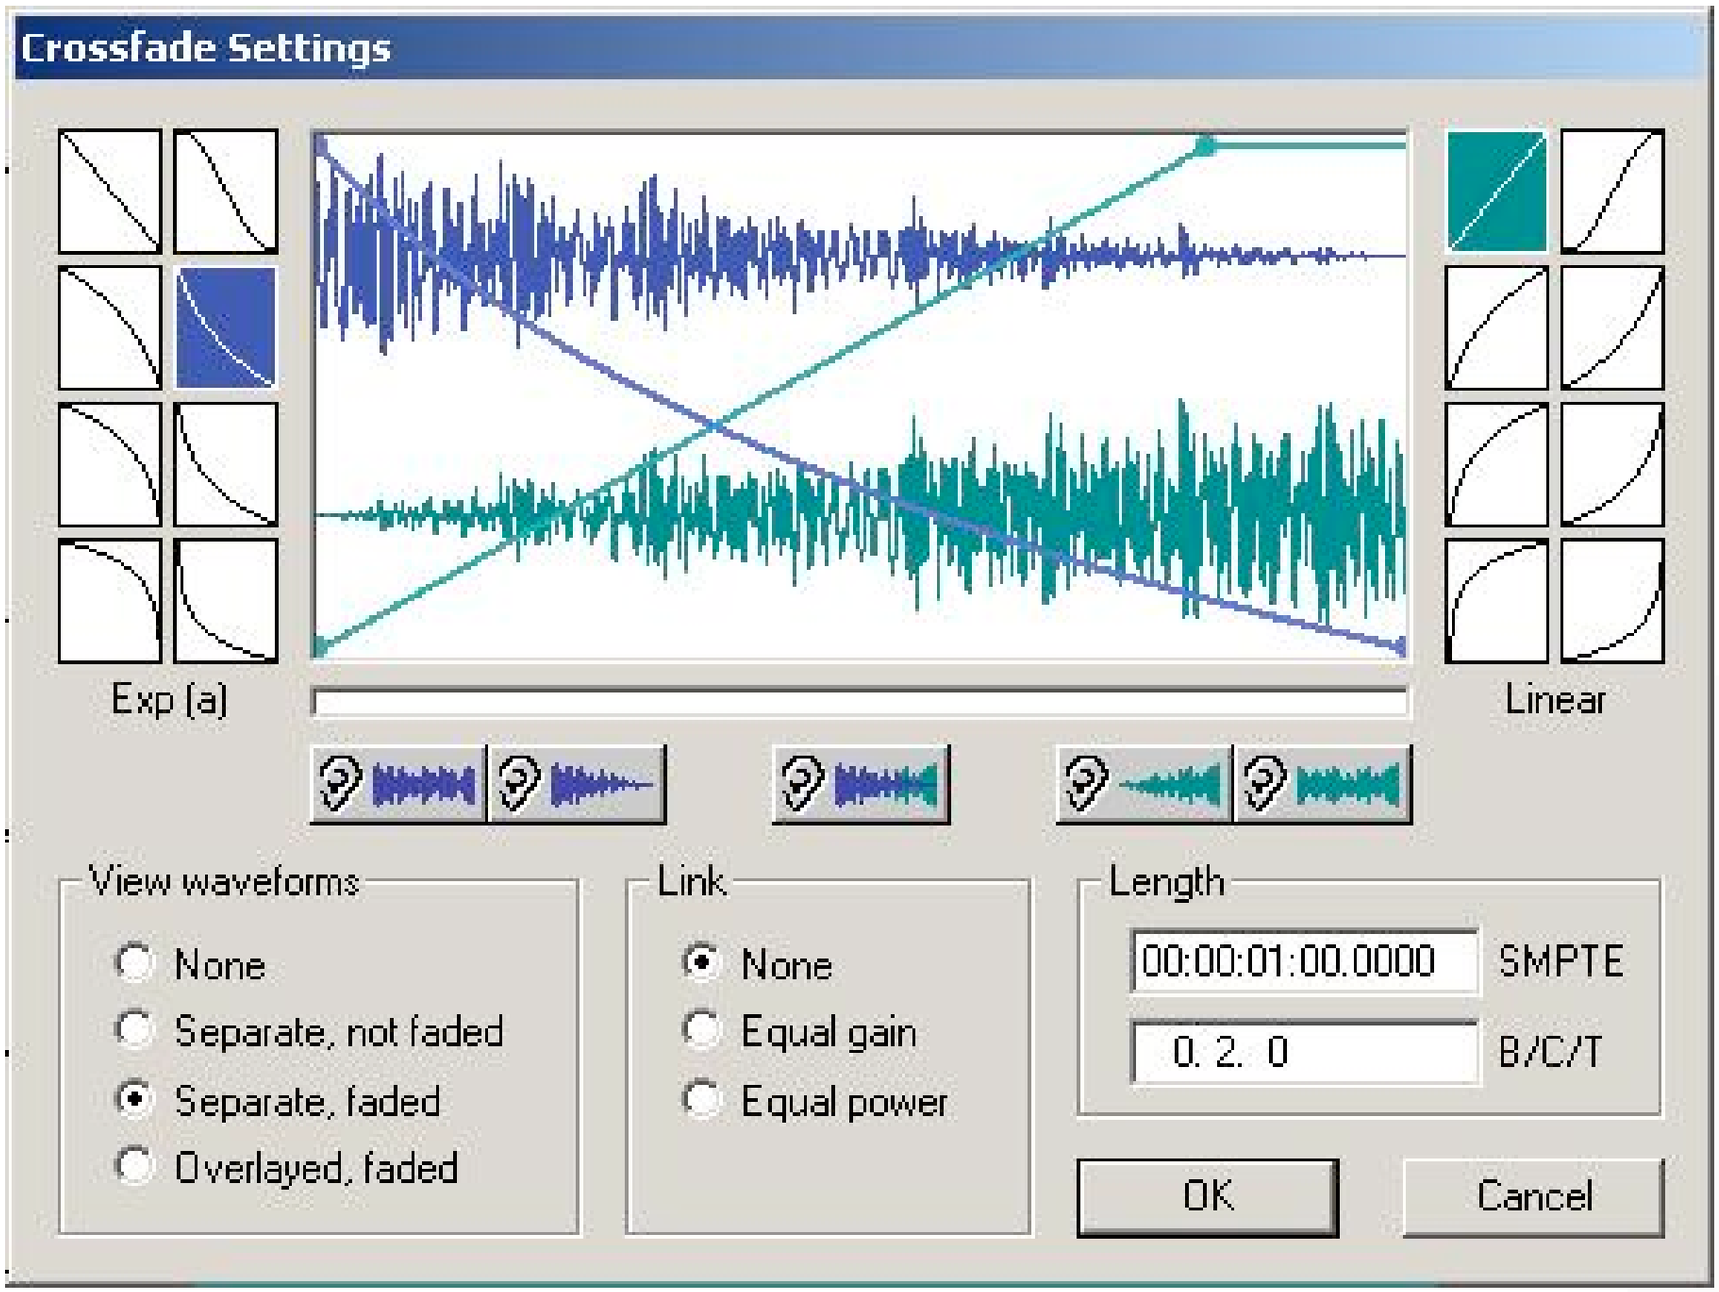

Each audio waveform multiplied by drawn gain contour

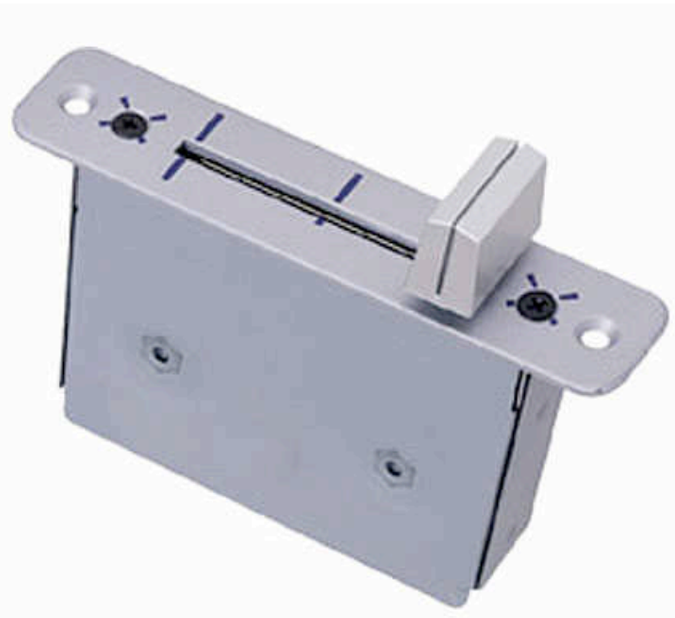

Shapes of fading functions yield different types of transitions.

Fusion Metrics

Fusion ...

Fusion metric. **The end of A and the start of B have different timbres (example: A is the "ta" onset of a trumpet, B is a sustained sound). We are looking for perceptual fusion across the splice.**

Why do we splice dissimilar timbres?

Because the transient at the start of a sounds forms a key part of the sound's identity to the listener.

History: Roland D-50, released in 1987

It combined a ROM of of short (100ms) samples of transients with a conventional synthesis engine for sustained sounds.

Very successful. Marked the end of FM synthesis era.

ROM samples (listen):

- Lips
- Pizz
- Steam

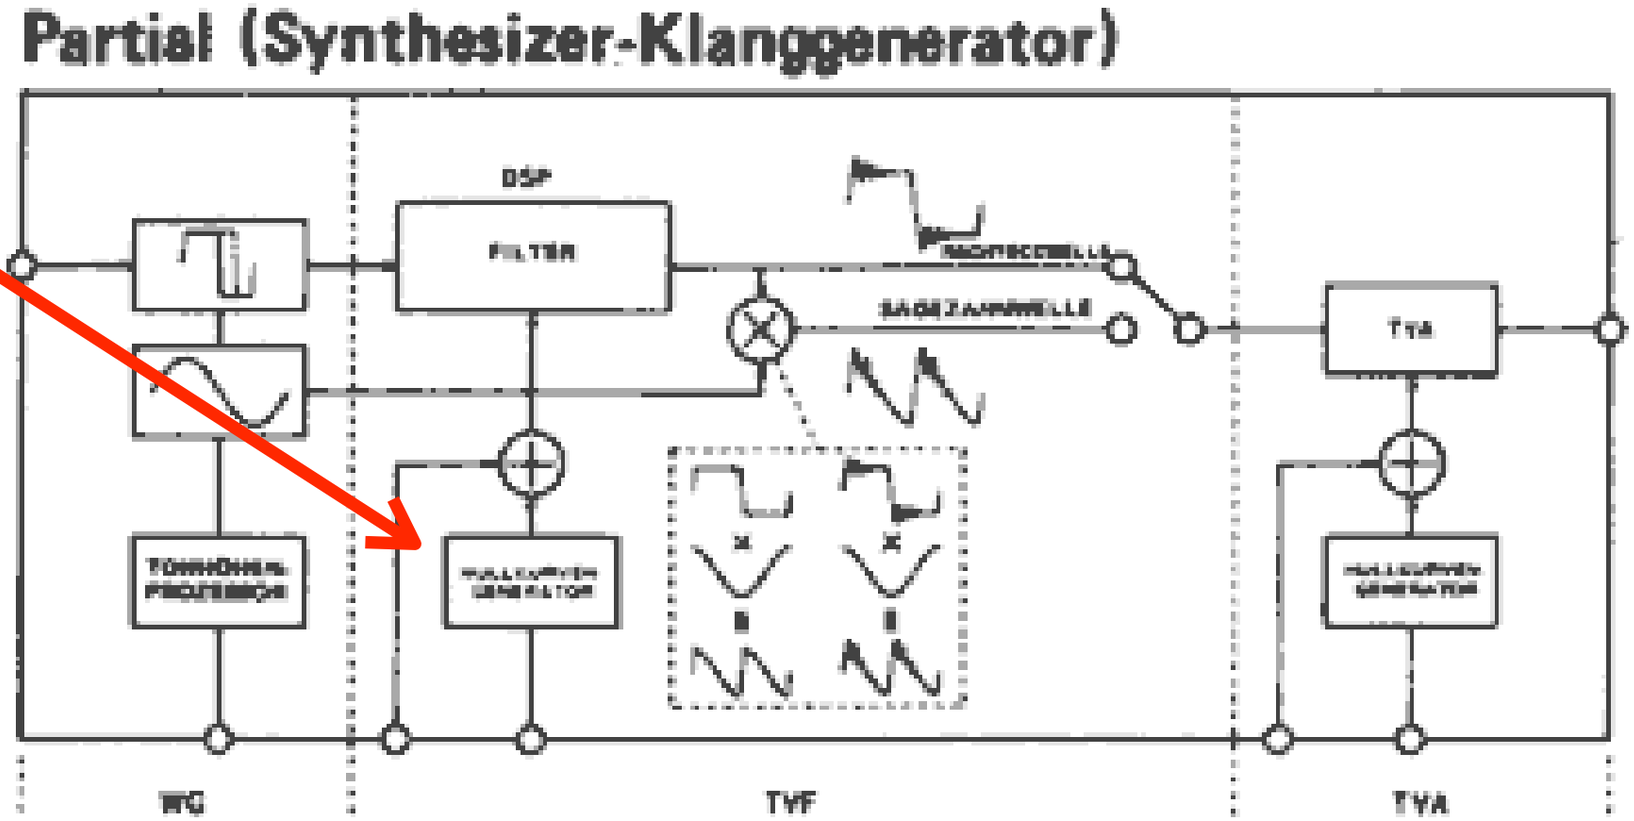

D50 patches using attack snippets

Acoustic Bass

Horn

Flute

Orinoco Flow

Why does the D-50 fuse so well?

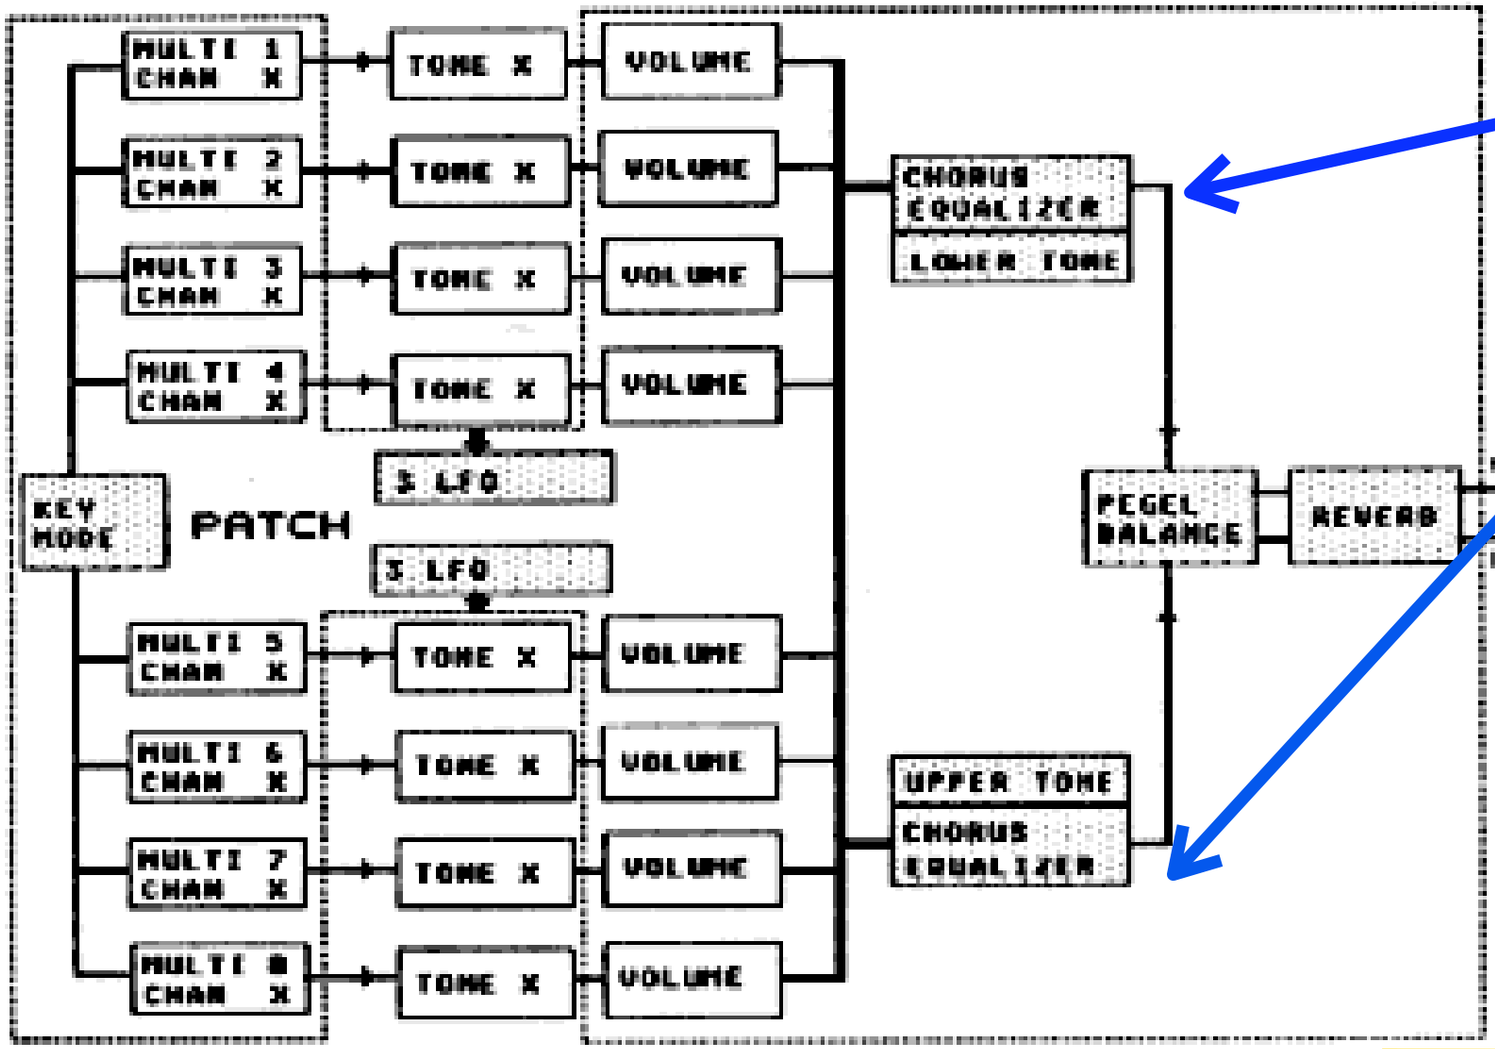

Built in chorus unit used to gloss over the transition.

But, sometimes this isn't the sound you want ...

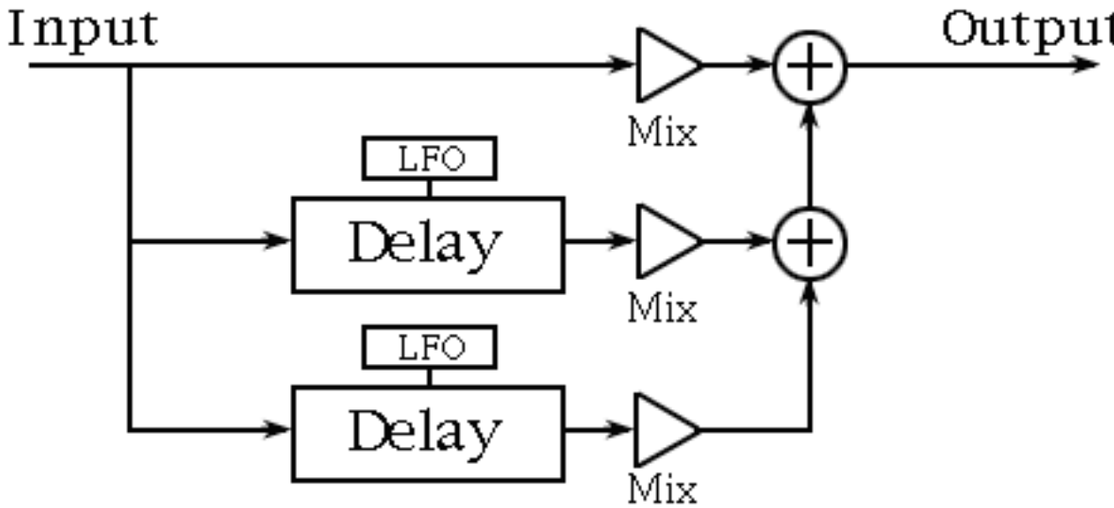

Another way to fuse onset snippets

SYNTHESIZING TRUMPET PERFORMANCES

Istvan Derenyi and Roger B. Dannenberg

School of Computer Science, Carnegie Mellon University
Pittsburgh, PA 15213, USA
{derenyi, rbd}@cs.cmu.edu

- **Each scale note has a trumpet onset sample.**
- **Measure the amplitude and phases of trumpet harmonics at the end of onset sample.**
- **To begin the sustained sound, a waveform is calculated whose phases and amplitudes match the onset.**
- **Over 50 ms, interpolate to the desired amplitude spectrum of the sustained sound.**

Demo of Dannenberg system ...

A real player

System with sampled attacks.

System without sampled attacks.

Resynthesis: Using this spectral approach to connect two samples, instead of connecting a sample to a synthesis algorithm (also known as spectral morphing).

Next: Eric Lindemann

SYNFUL

ORCHESTRA