Music 209
Advanced Topics in Computer Music
Lecture 4 – Time Warping

2006-2-9

Professor David Wessel (with John Lazzaro)

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

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

Today, we focus on time warping techniques to modify the duration of candidates.

Select candidate samples from db

Any good matches?

N

Modify a candidate to be good enough

#3(mod)

Do the splice

#3(mod)

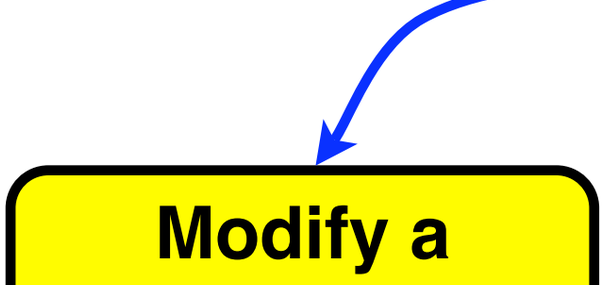

Y

Choose best match

#2

Do the splice

#2

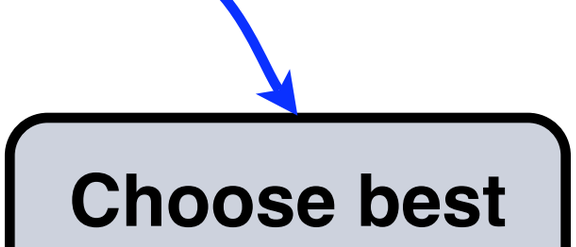

Candidates

#1

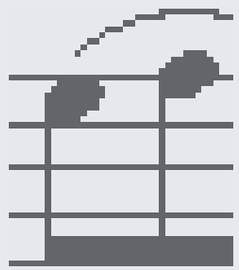

#2

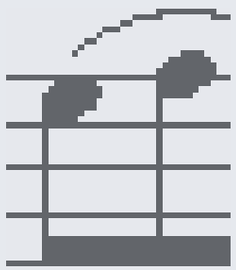

#3

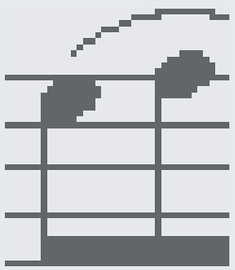

Topics for today ...

- **Time-warping: change the duration of musical audio in a “transparent” way.**
- **Finishing material from earlier lectures.**
- **Project ideas on parade.**

define: Note-level Time Warping

Time Shrink

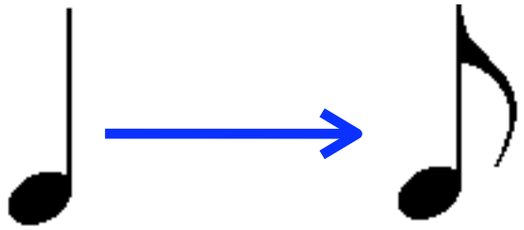

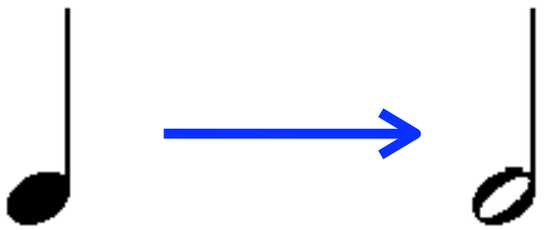

Time Stretch

- **Length of attack transient unchanged.**
  Time warp only affects sustained region.

- **Local temporal properties of sustained region unchanged (example: vibrato speed)**

- **Long-range properties of sustained region stretch or shrink (example: crescendos).**

Concept: Like a player would lengthen/shorten note.

define: Phrase-level Time Warping

original: T = 2s

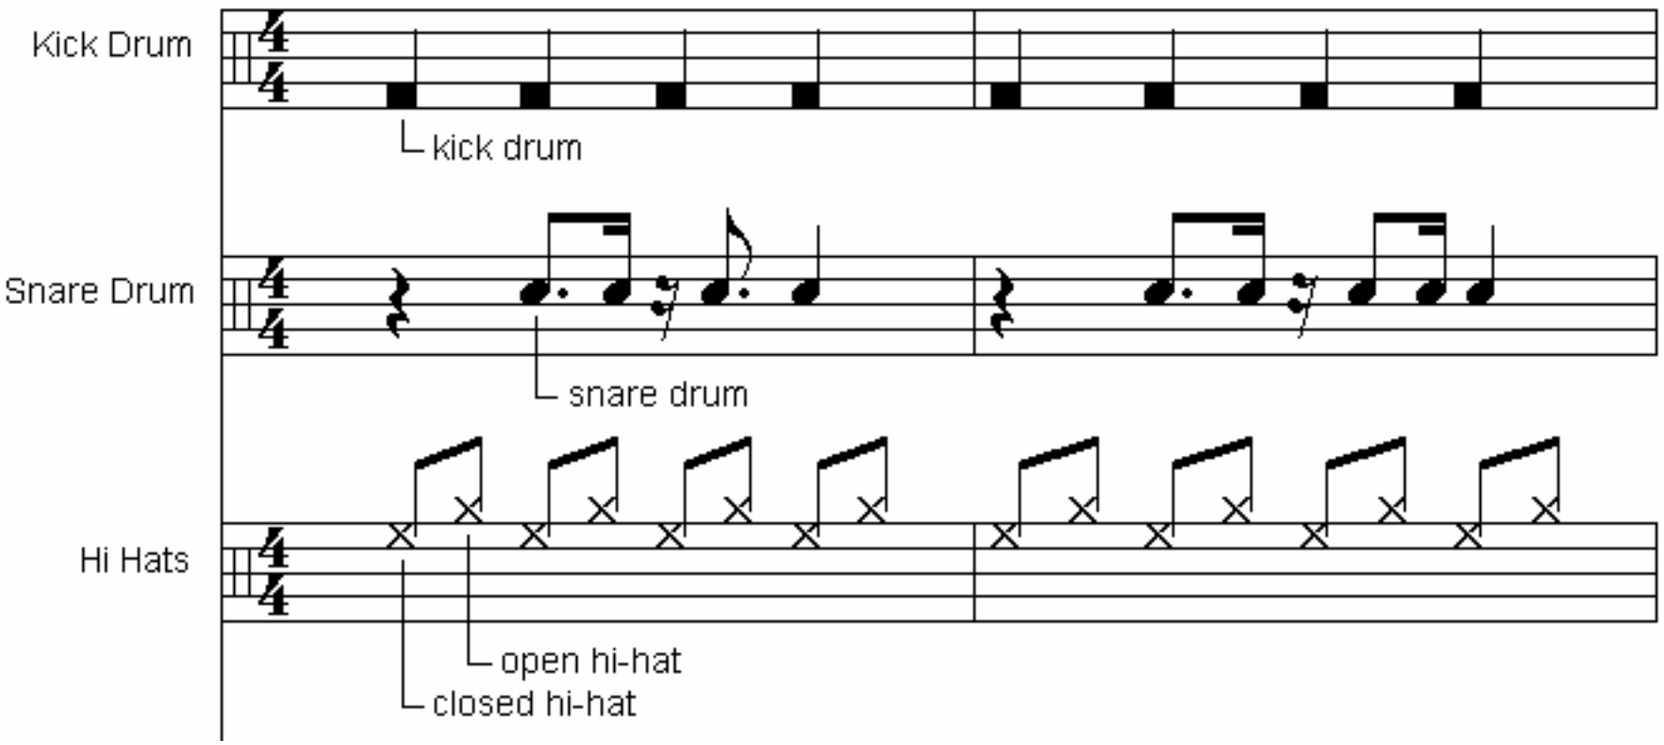

Time Shrink T = 1.5s

Time Stretch T = 2.5s

- For **percussive** sounds, shrink or stretch by **adding** or **removing rests** (silence).
- For **sustained** sounds, scale note **durations** and **rests** to match the score.

Concept: Like a robot player would change tempo.

THIS SEEMS LIKE AN AI-COMPLETE PROBLEM ...

Different algorithms sound better on some sound types and worse on others, because they make different assumptions about the source material.

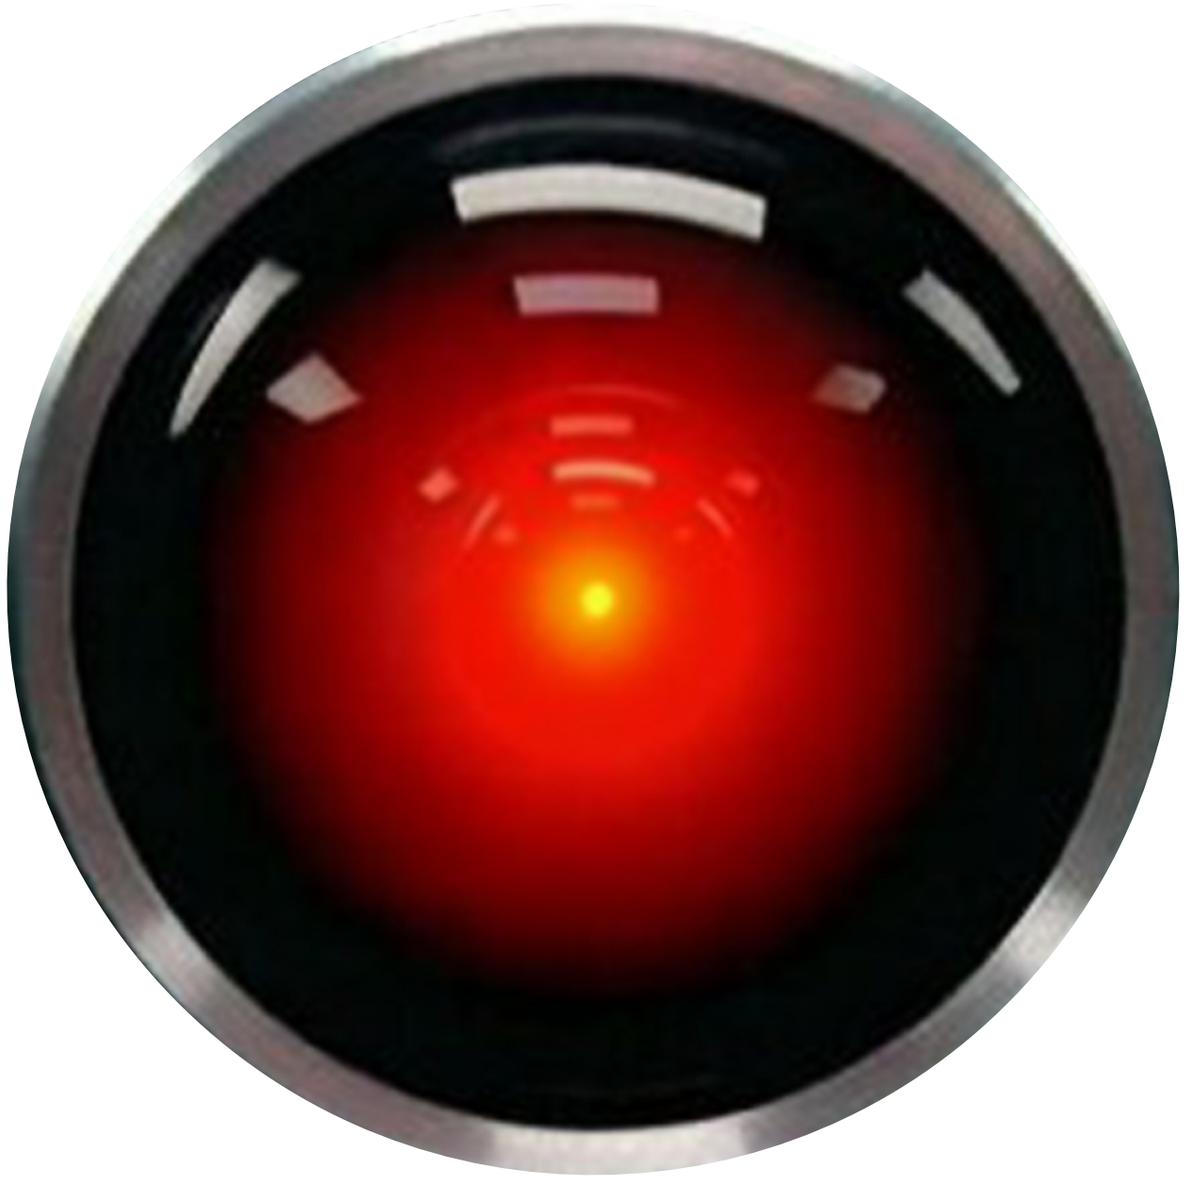

And so, we cover many approaches to this problem in the lecture today ...

Drum Phrase Algorithms

We begin with two ways to sidestep the time-shifting problem ...

Sidestep #1: DrumCore (Submersible Music)

Play
Play

Hire drummers famous in different genres.

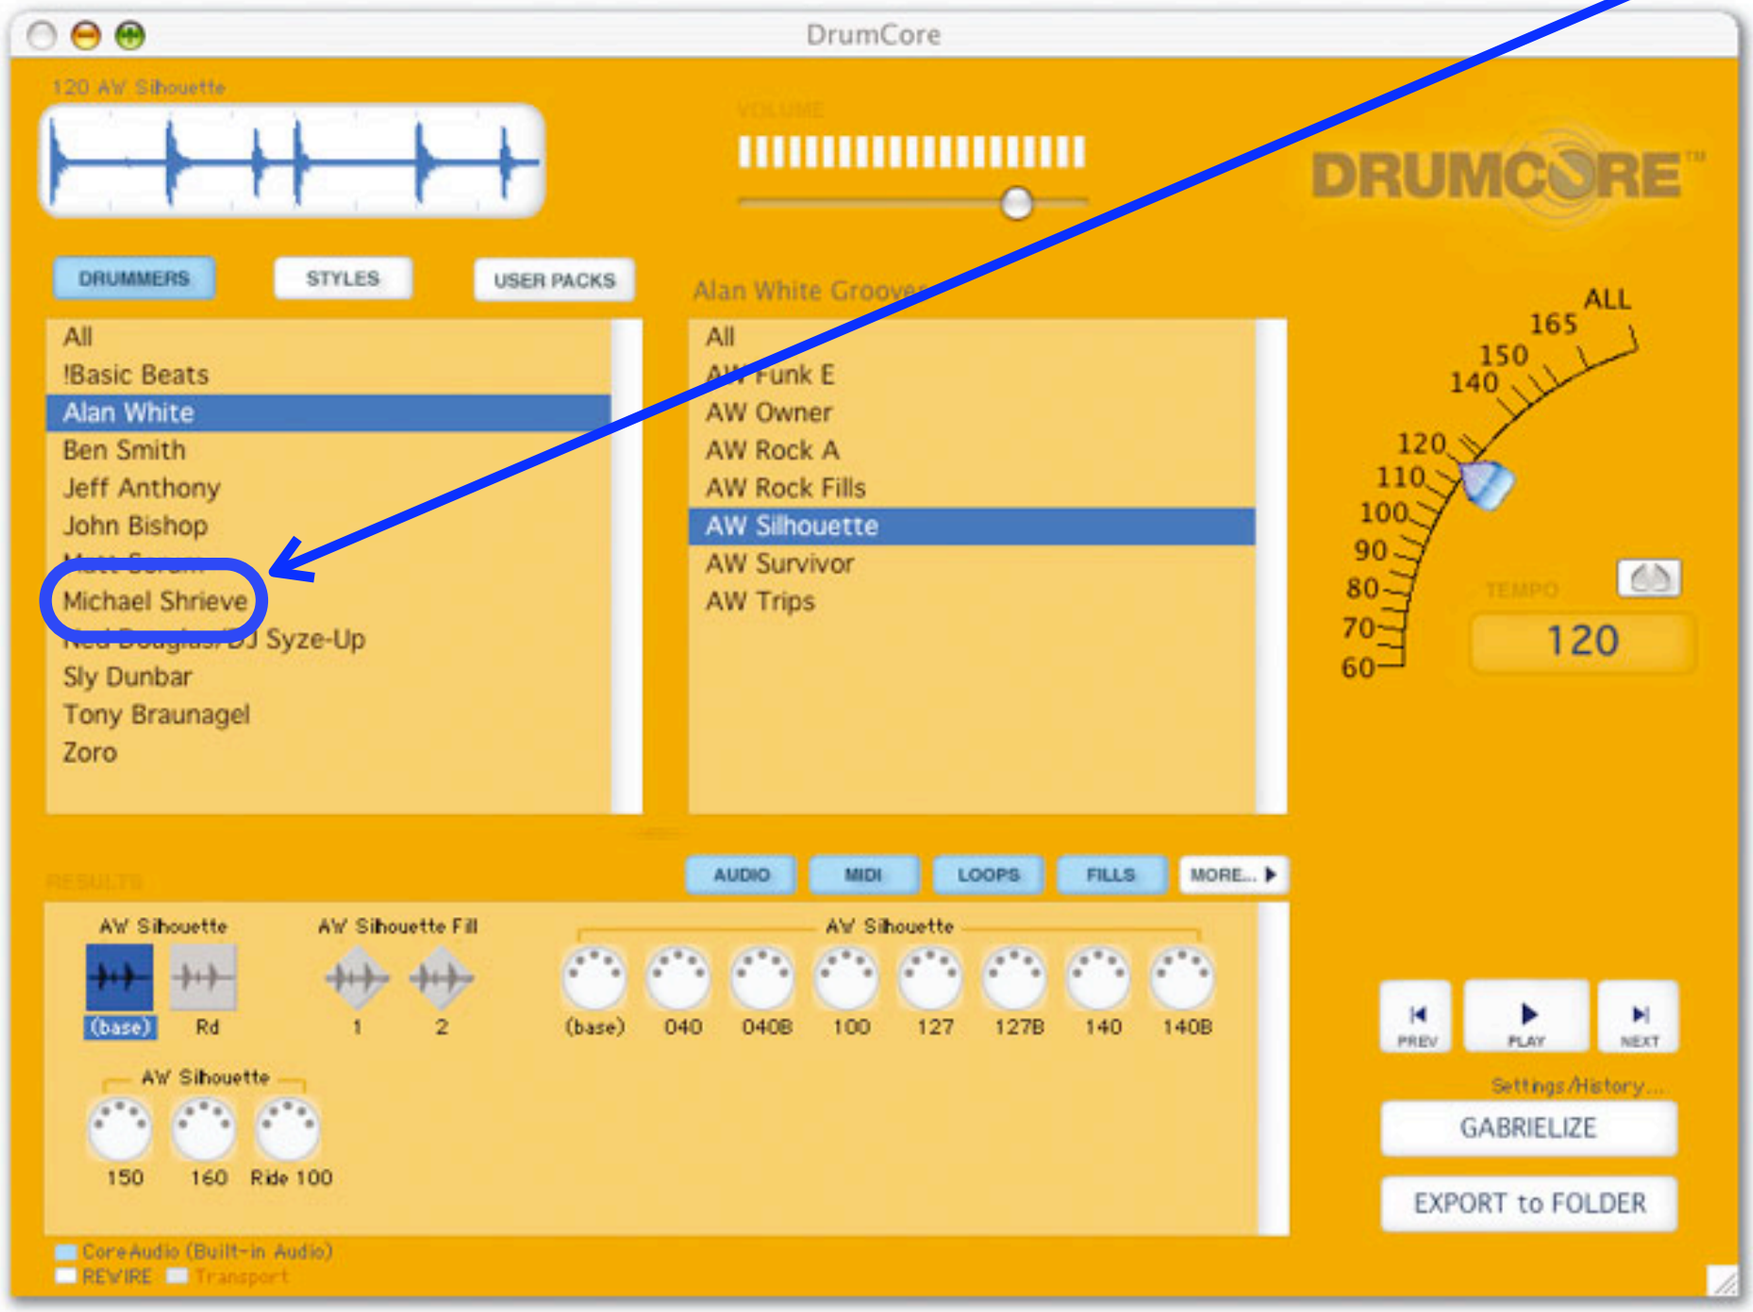

Record them playing the same beat at many tempos.

Recall our goal: Like a robot player would change tempo.

Drumcore doesn't do this! A drummer plays the same beat differently at 60 bpm and 165 bpm ... hmm ...

Sidestep #2: Sample-based drum machines.

Sample data holds a single drum hit sound.

Change tempo by changing the rate the drum sounds trigger.

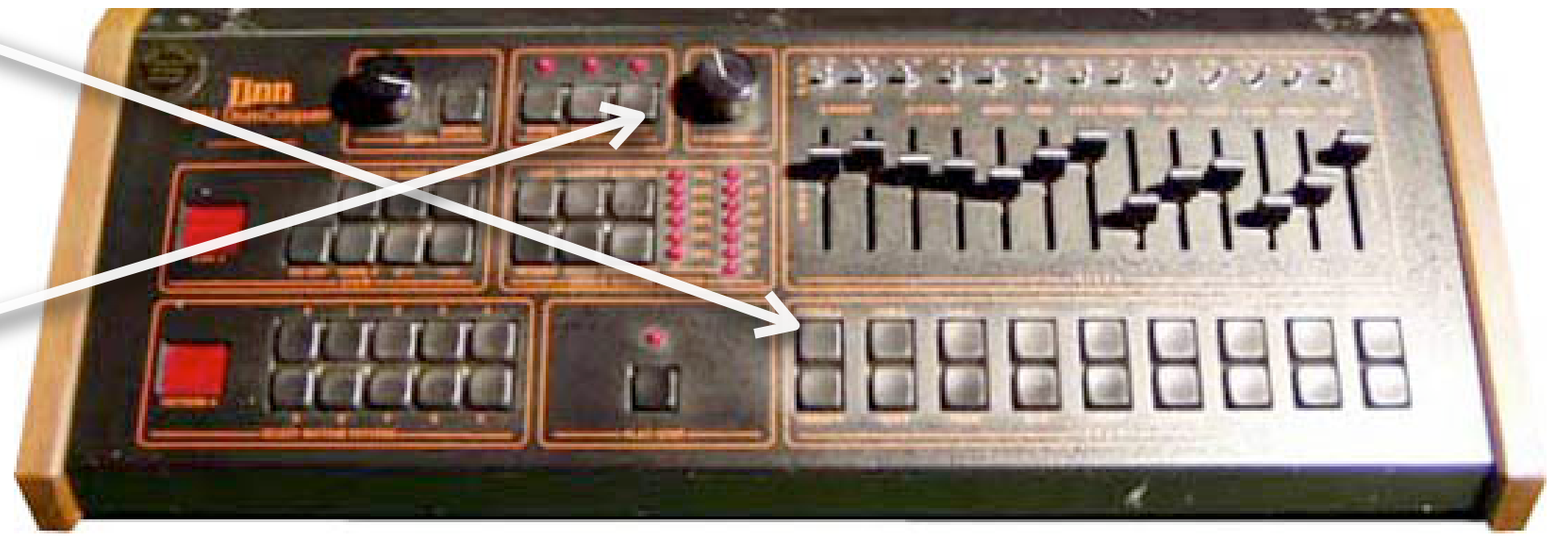

Recall our goal: Time-warp audio recording of a drum pattern like a robot player would change tempo.

Can't we directly apply the drum machine model?

Problem #1: Extracting clean isolated drum hit sounds for every drum sound in the recording.

Problem #2: Extracting the "trigger pattern" for each drum sound without error is a hard problem.

However, a variant of this idea works quite well …

Beat Slicing

Invented as a manual process ...

Later automated by Propellorhead Software's ReCycle, and extended by Spectrasonics.

File types:

REX (Recycle),
Groove Control (Spectrasonics)

Basic idea: Use transient detection to divide an audio recording into shorter contiguous slices.

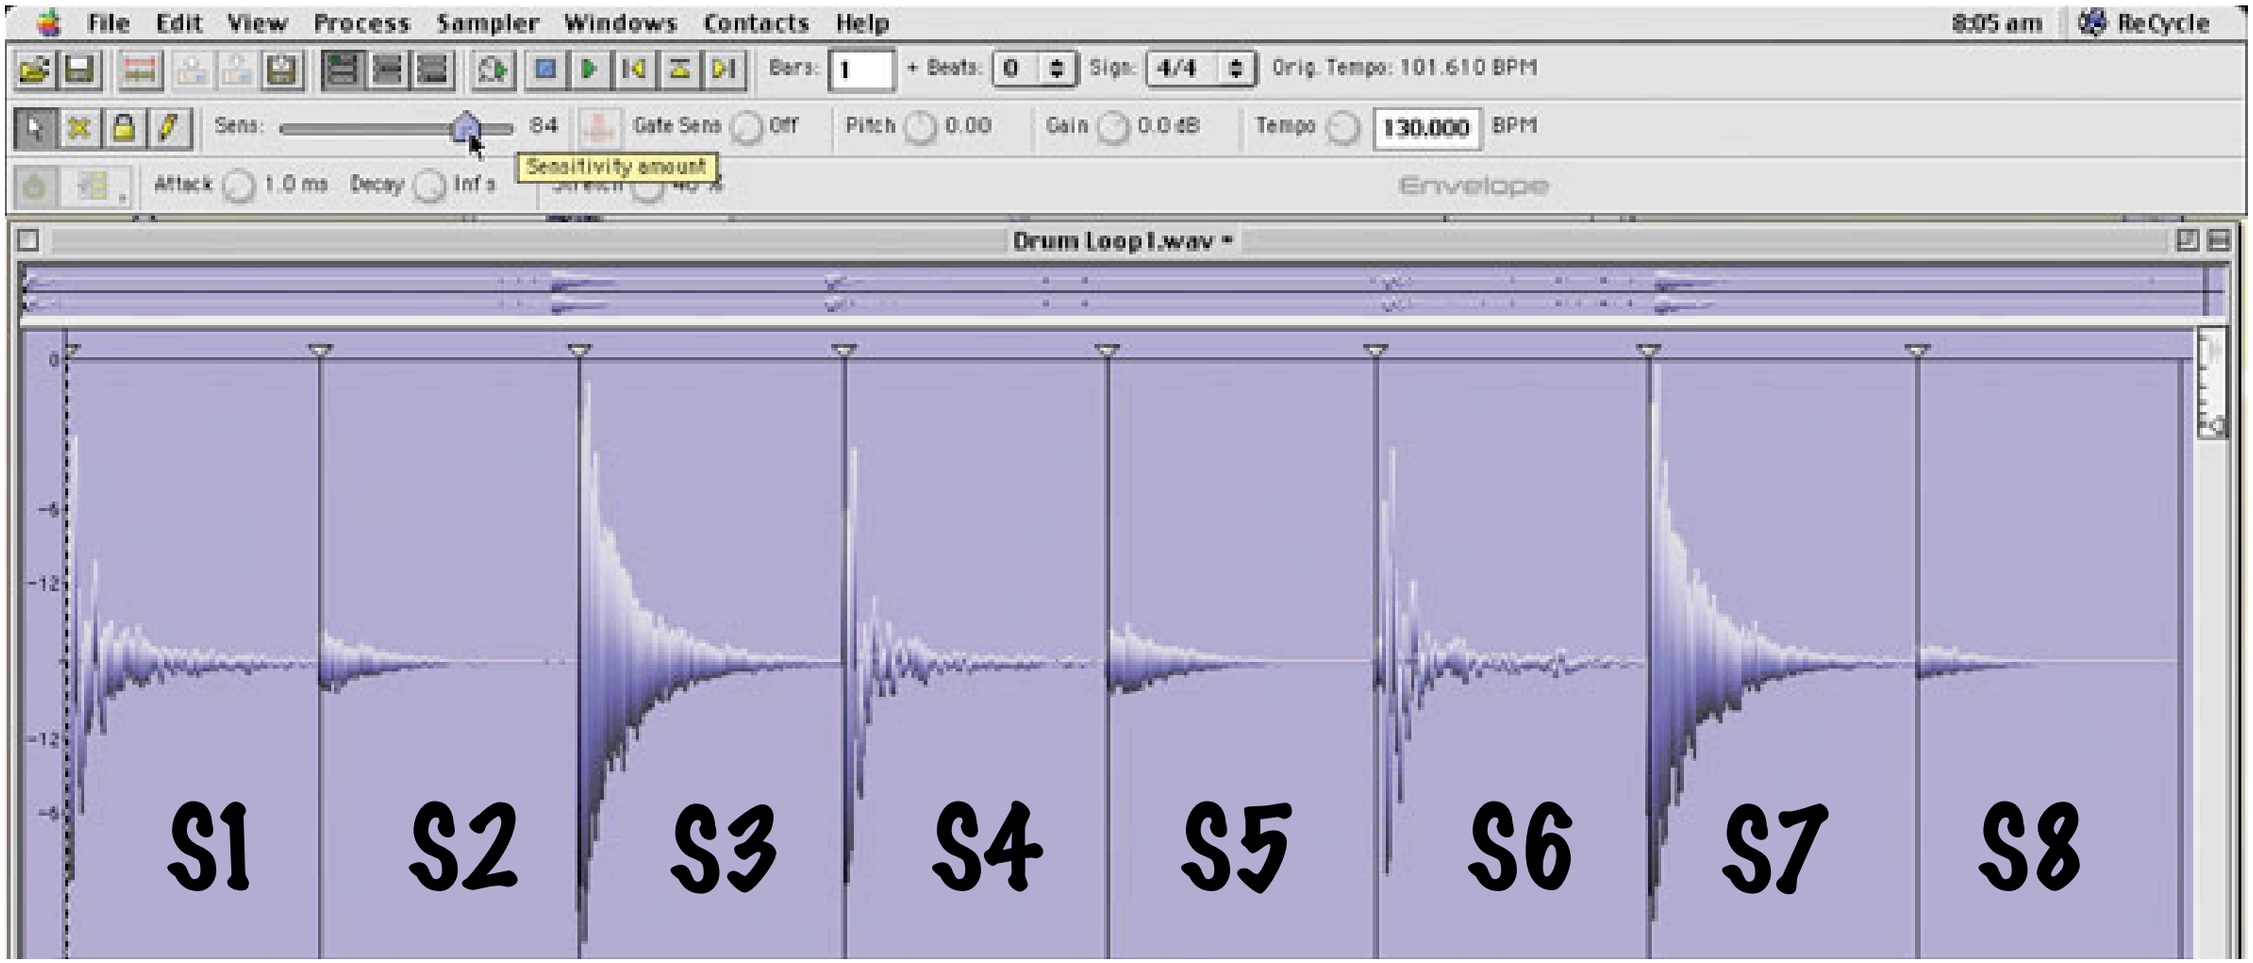

Assign each numbered sample to its own MIDI note number (S1 is triggered by C#4, S2 is triggered by D4, S3 is triggered by D#4, ...).

EXAMPLE BY SIMON PRICE, IN JUNE 2004 SOUND ON SOUND MAGAZINE

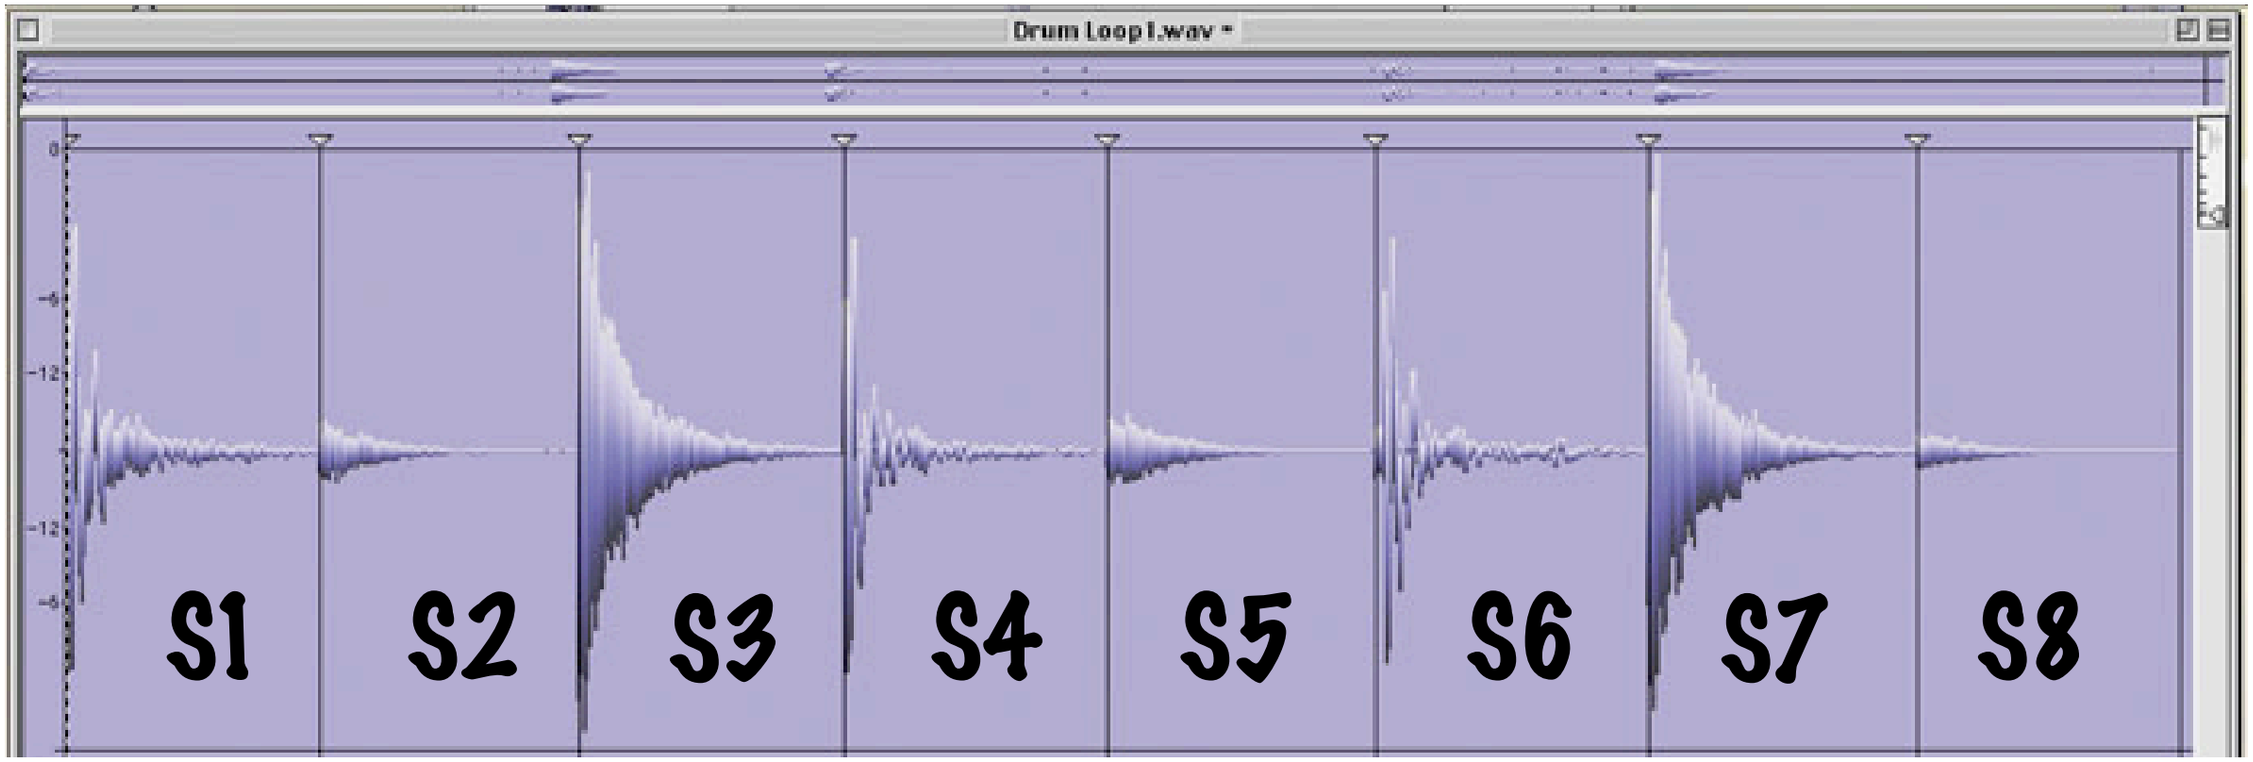

Create a MIDI performance of a chromatic scale, whose note timing trigger each sample at the perfect time to recreate the original audio.

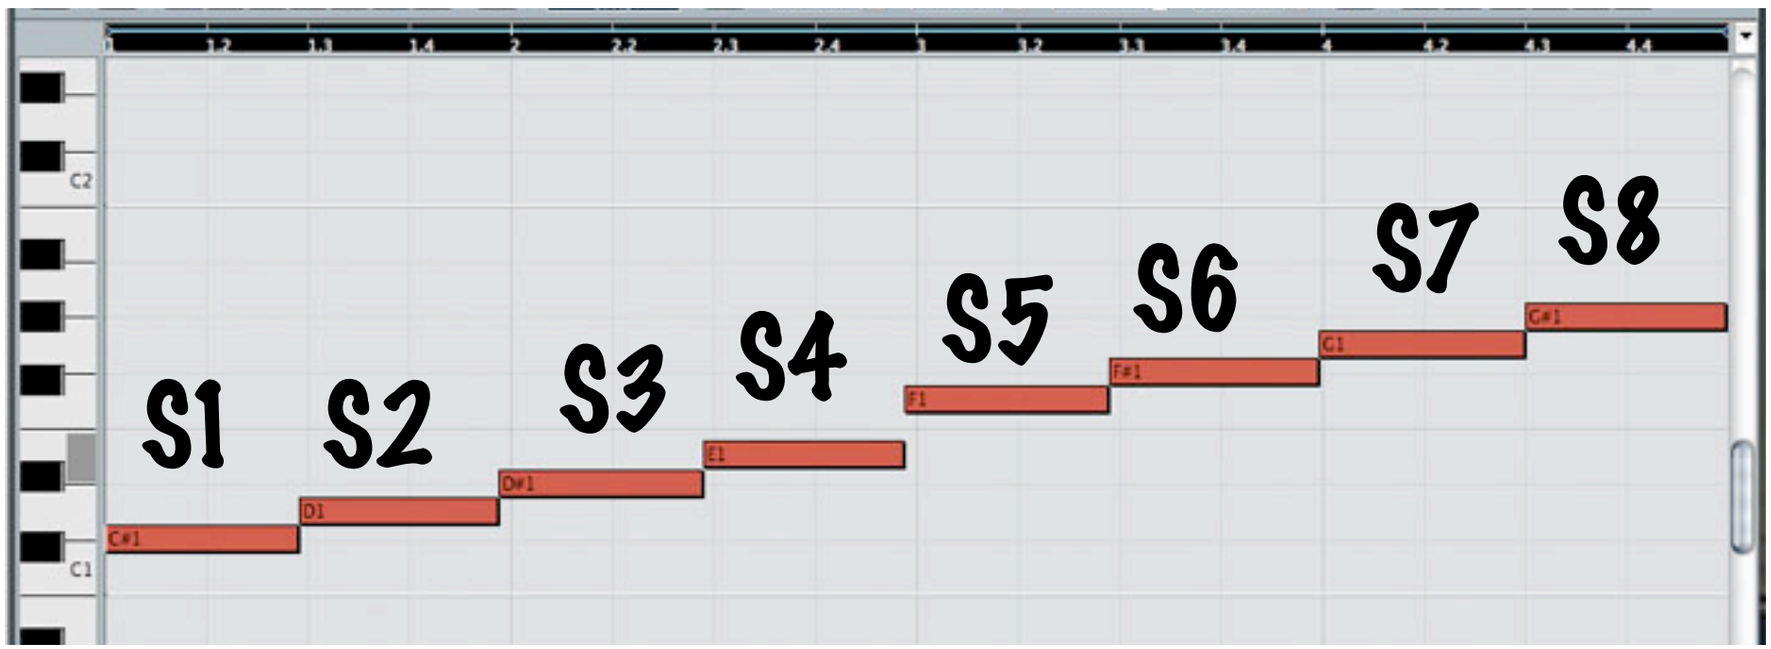

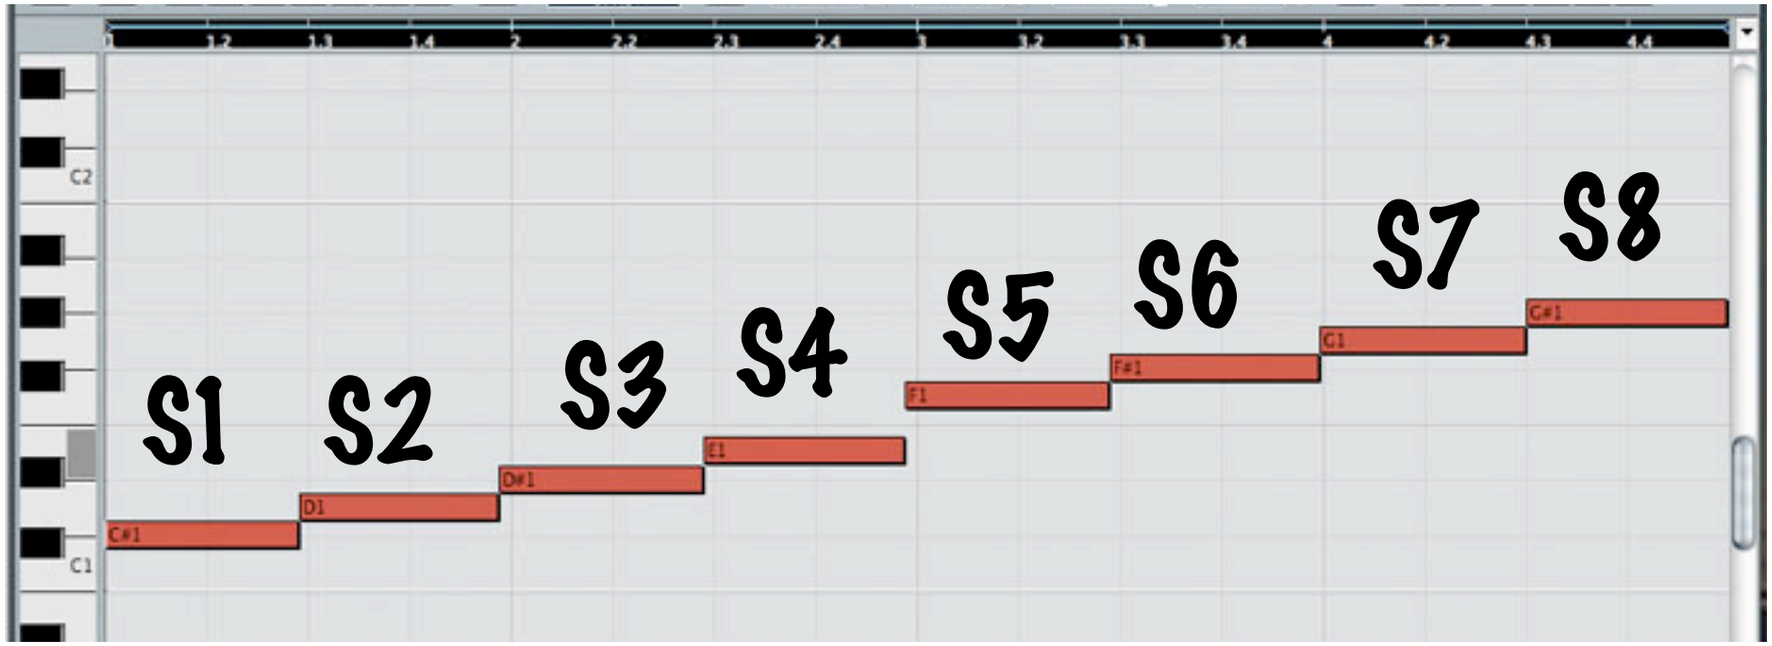

To time-stretch, slow down MIDI tempo:

Extra time introduced as silent gaps ...

EXAMPLE BY LEN SASSO, IN OCT 2003 ELECTRONIC MUSICIAN MAGAZINE.

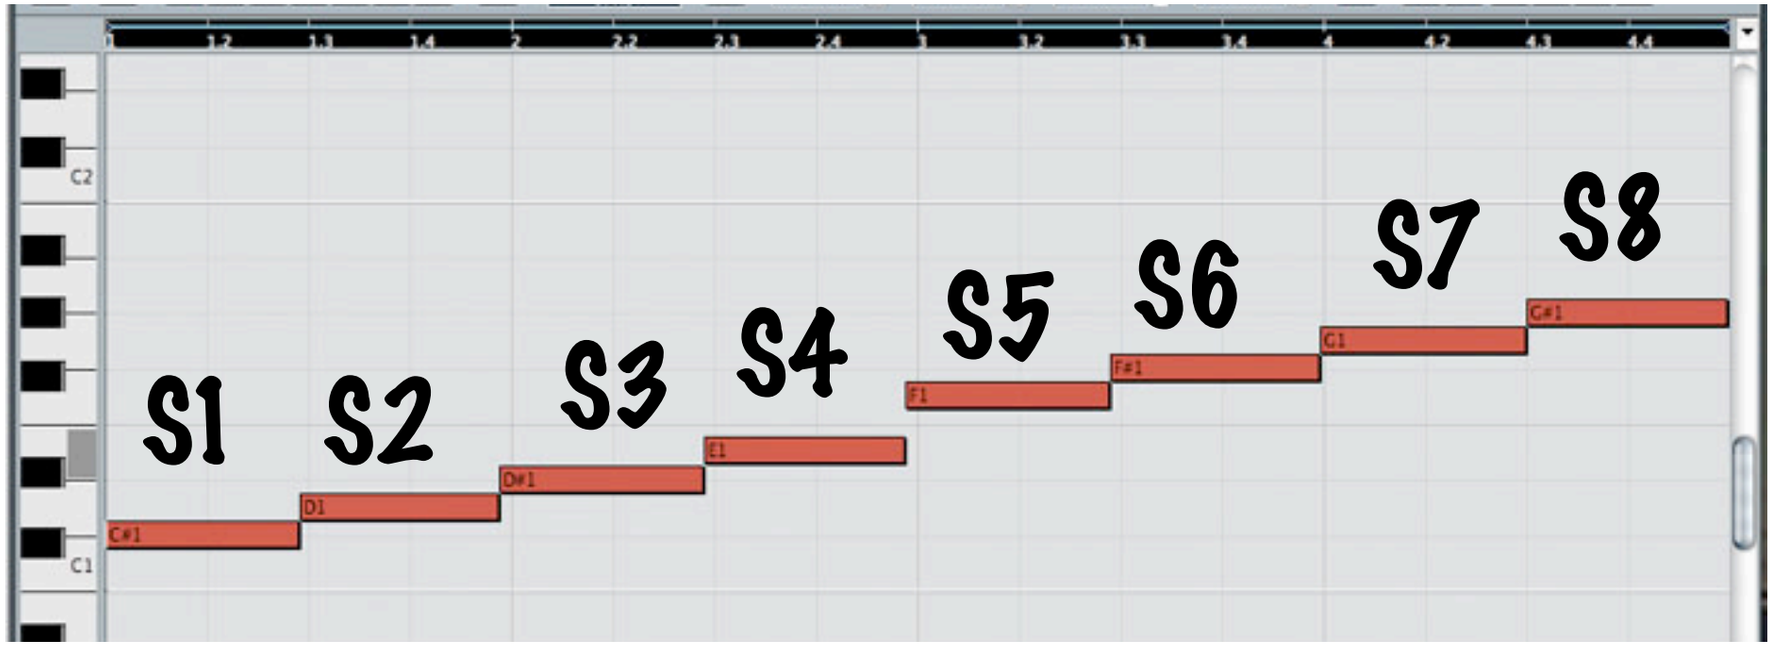

To time-shrink, speed up MIDI tempo:

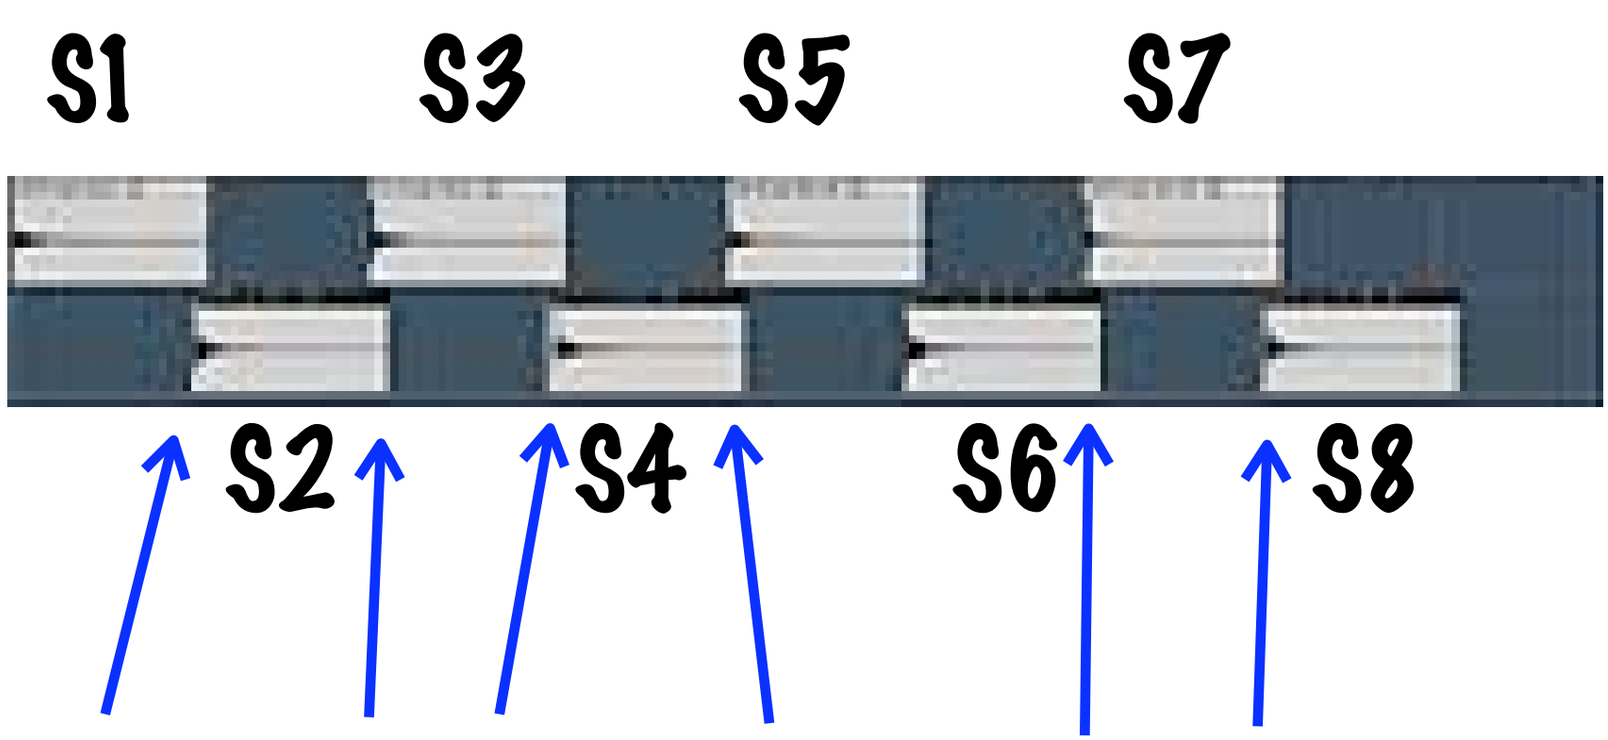

Overlaps reduce the total playing time ...

EXAMPLE BY LEN SASSO, IN OCT 2003 ELECTRONIC MUSICIAN MAGAZINE.

Artifacts ...

Good news: attacks are artifact free. The most important part of a percussion sound.

Artifact: Silent gap.

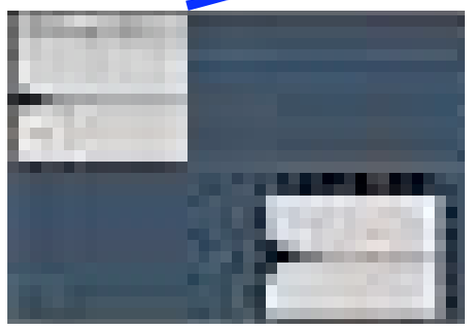

Artifact: Tail overlap.

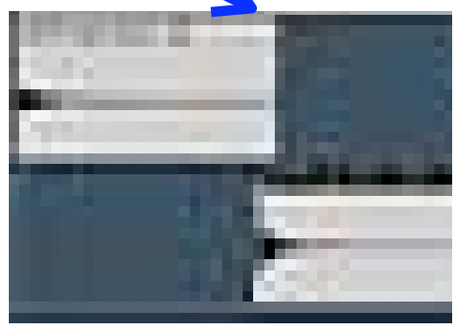

- **One approach: Apply envelope to each slice to fade it to silence before gap or overlap.**
- **Another approach for gaps: loop the end of the tail to extend it through the gap.**

Another option: sample substitution

Gap before this HH will be noticeable because of kick tail.

Gap before this HH might sound better.

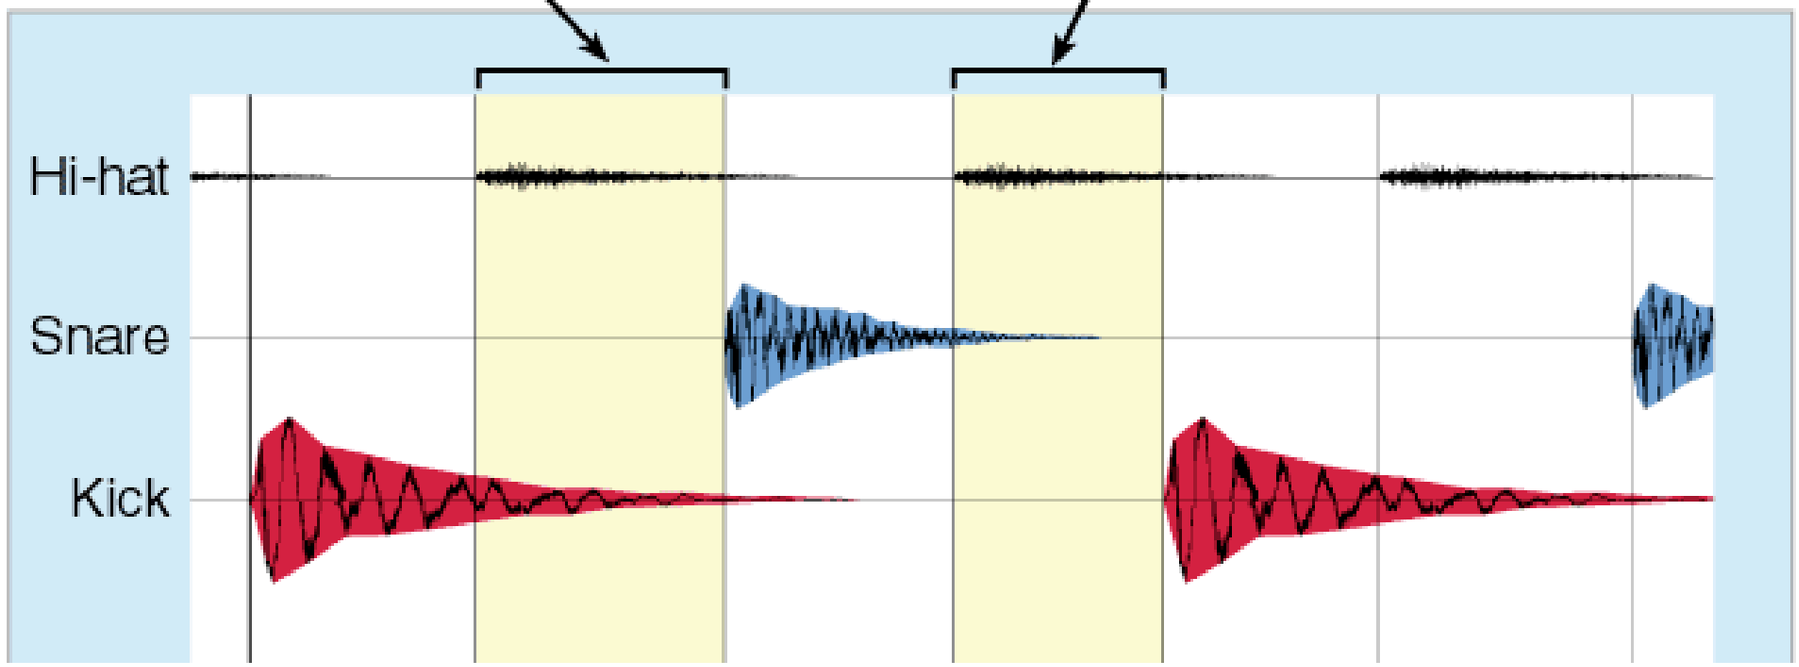

So, use S4 sample to replace S2.

How does beat slicing sound?

- **Before:** 5 audio tracks (hi-hat, bass drum, polysynth, electric piano, bass) at 5 different beats per minute (120, 130, 126, 135, 114).

  Play

- **After:** Each track beat sliced by using a transient detection algorithm and gap and tail artifact reduction (ReCycle + Phasmatic Pro)

  Play

- **A simpler approach:** Each track beat sliced, with slices created at 16th-note intervals (Ableton Live, 2003 version).

  Play

EXAMPLE BY LEN SASSO, IN OCT 2003 ELECTRONIC MUSICIAN MAGAZINE.

Sample substitution: Artistic uses ...

S1 S2 S3 S4 S5

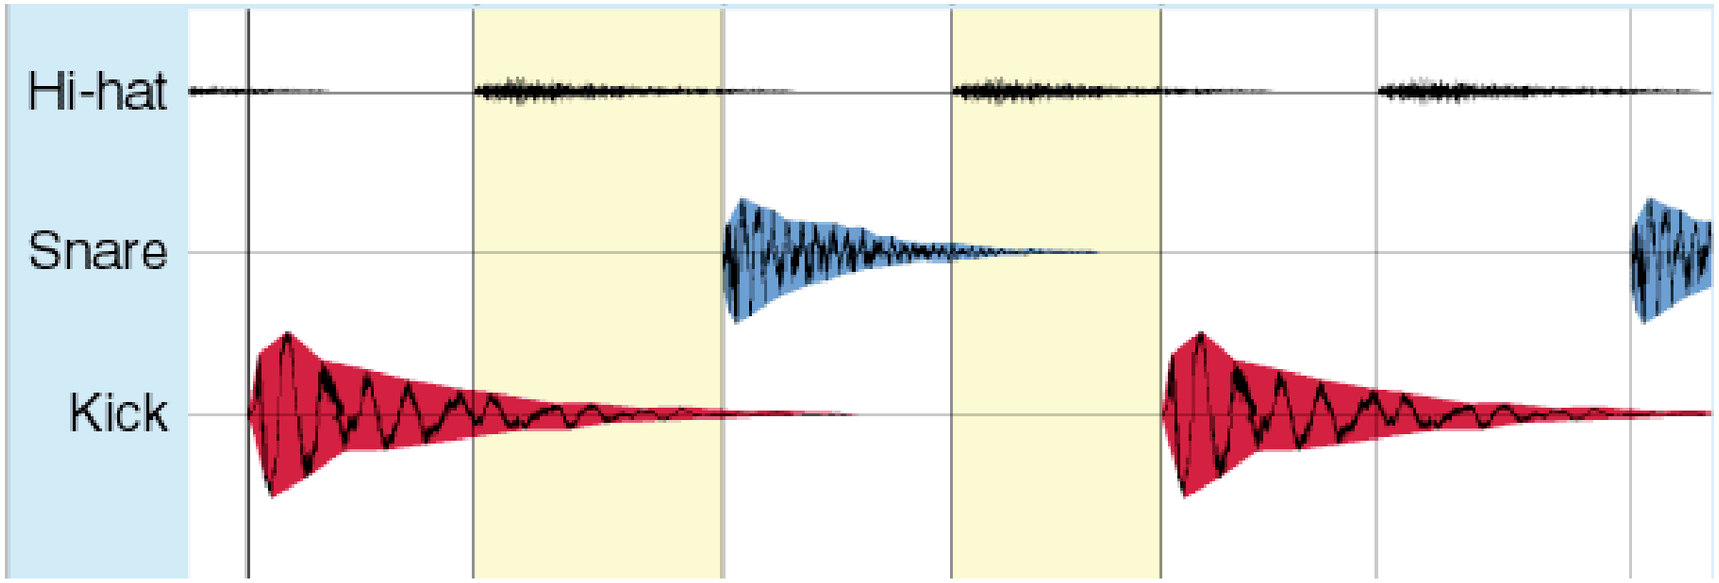

- **Original beat:** Play
- **Reworked beat:** Play
- **Client: “Where’s the chuffa-chuffa?”**

The artifacts become a feature ... Play

Pro-Tools Beat Detective ...

Record a live drummer (with many mics).
Beat slice performance to fix “feel” issues.

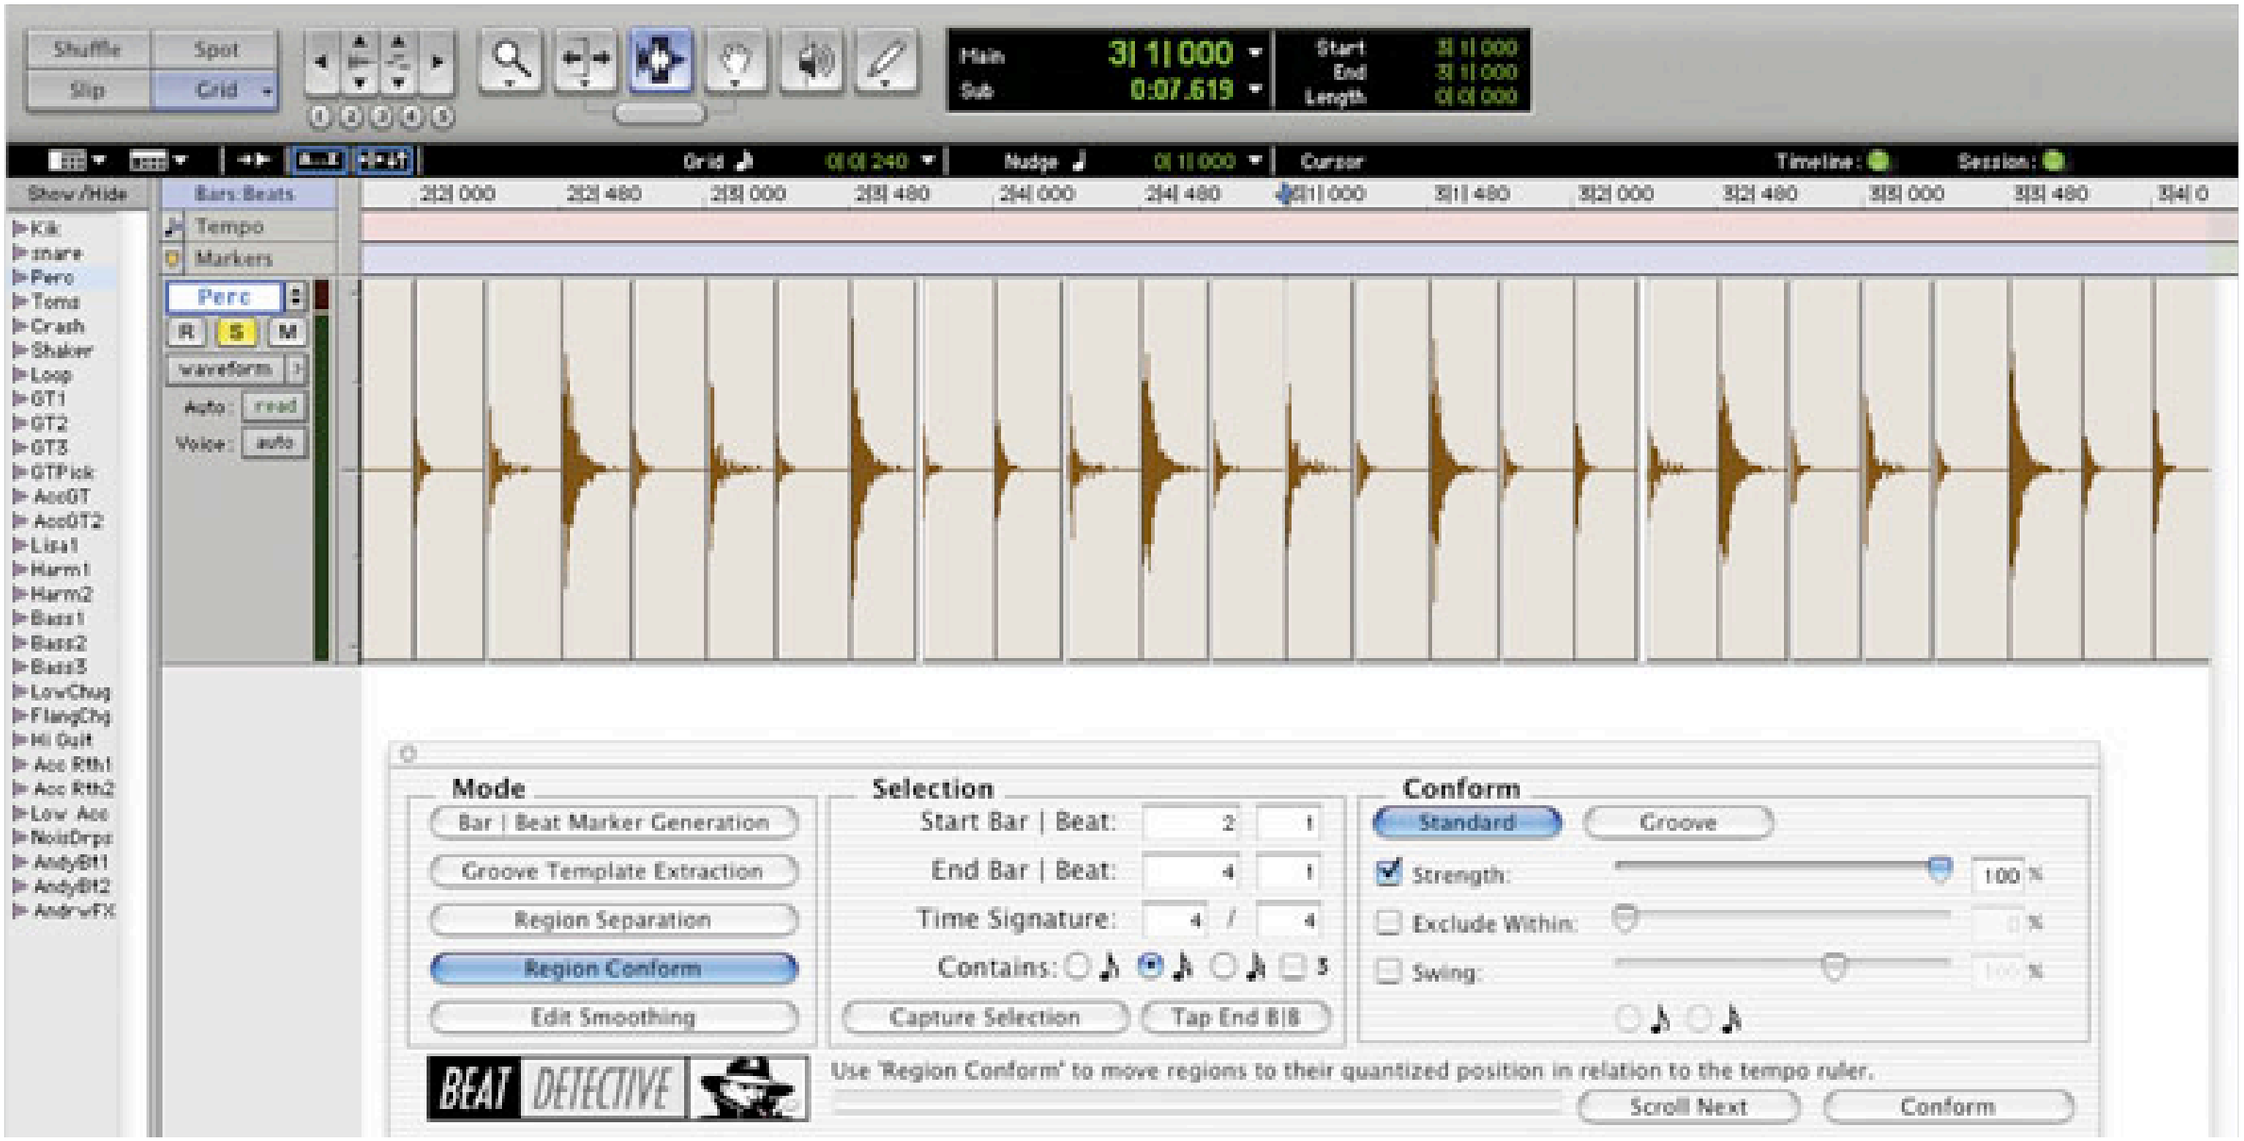

Young engineer: You need Beat Detective to cut a live drummer.

Old engineer: You only need Beat Detective if your drummer is Beat Defective.

Recall: Note-level Time Warping

We now look at time warping algorithms for this problem ...

Time Shrink

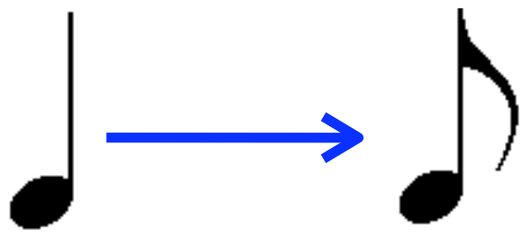

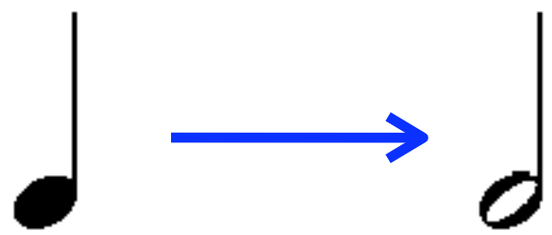

Time Stretch

- **Length of attack transient unchanged. Time warp only affects sustained region.**
- **Local temporal properties of sustained region unchanged (example: vibrato speed)**
- **Long-range properties of sustained region stretch or shrink (example: crescendos).**

Overlap-Add Methods

Commercially in use since the 1980s - Lexicon 2400.

Intro to Overlap and Add Methods ...

Step 1: Create fixed-length "sample blocks" $x_1(n)$, $x_2(n)$, $x_3(n)$ ... that overlap 50% in time.

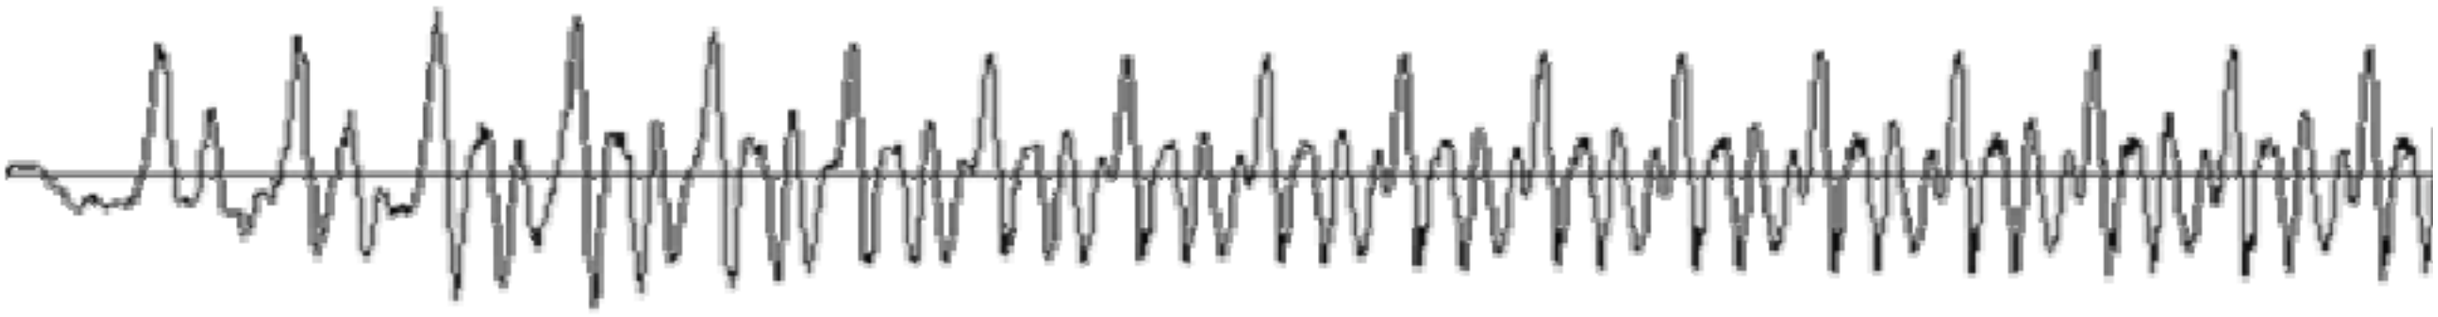

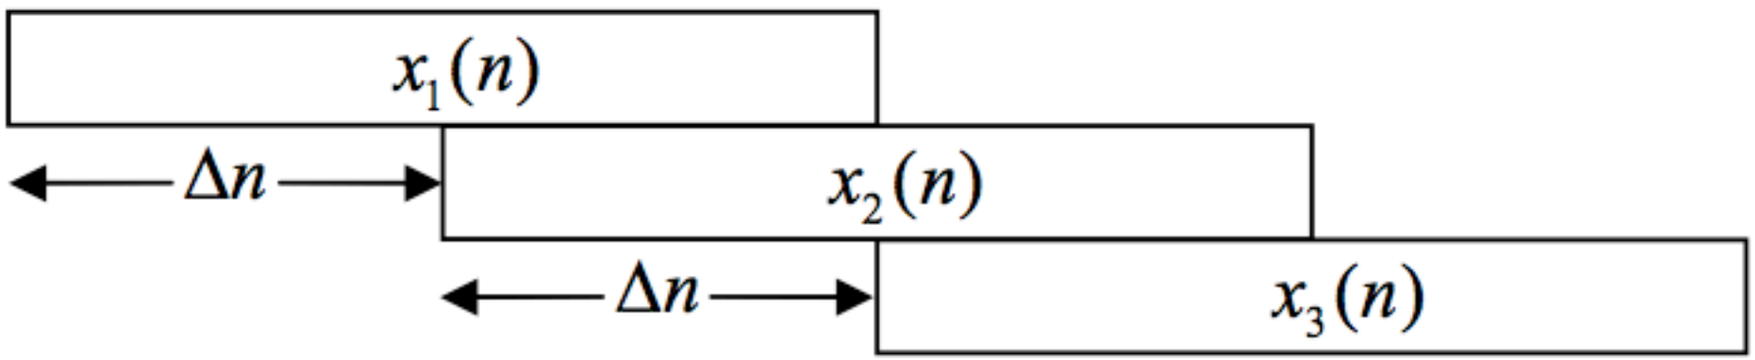

EXAMPLES FROM JORDI BONADA SANJAUME, BARCELONA

Step 2: Define α to be the time-scale factor.

$\alpha = 1$ denotes no scaling. $\alpha > 1$ is time-stretch.

$\alpha < 1$ is time-shrink. **Offset blocks by:** $\alpha \Delta n$

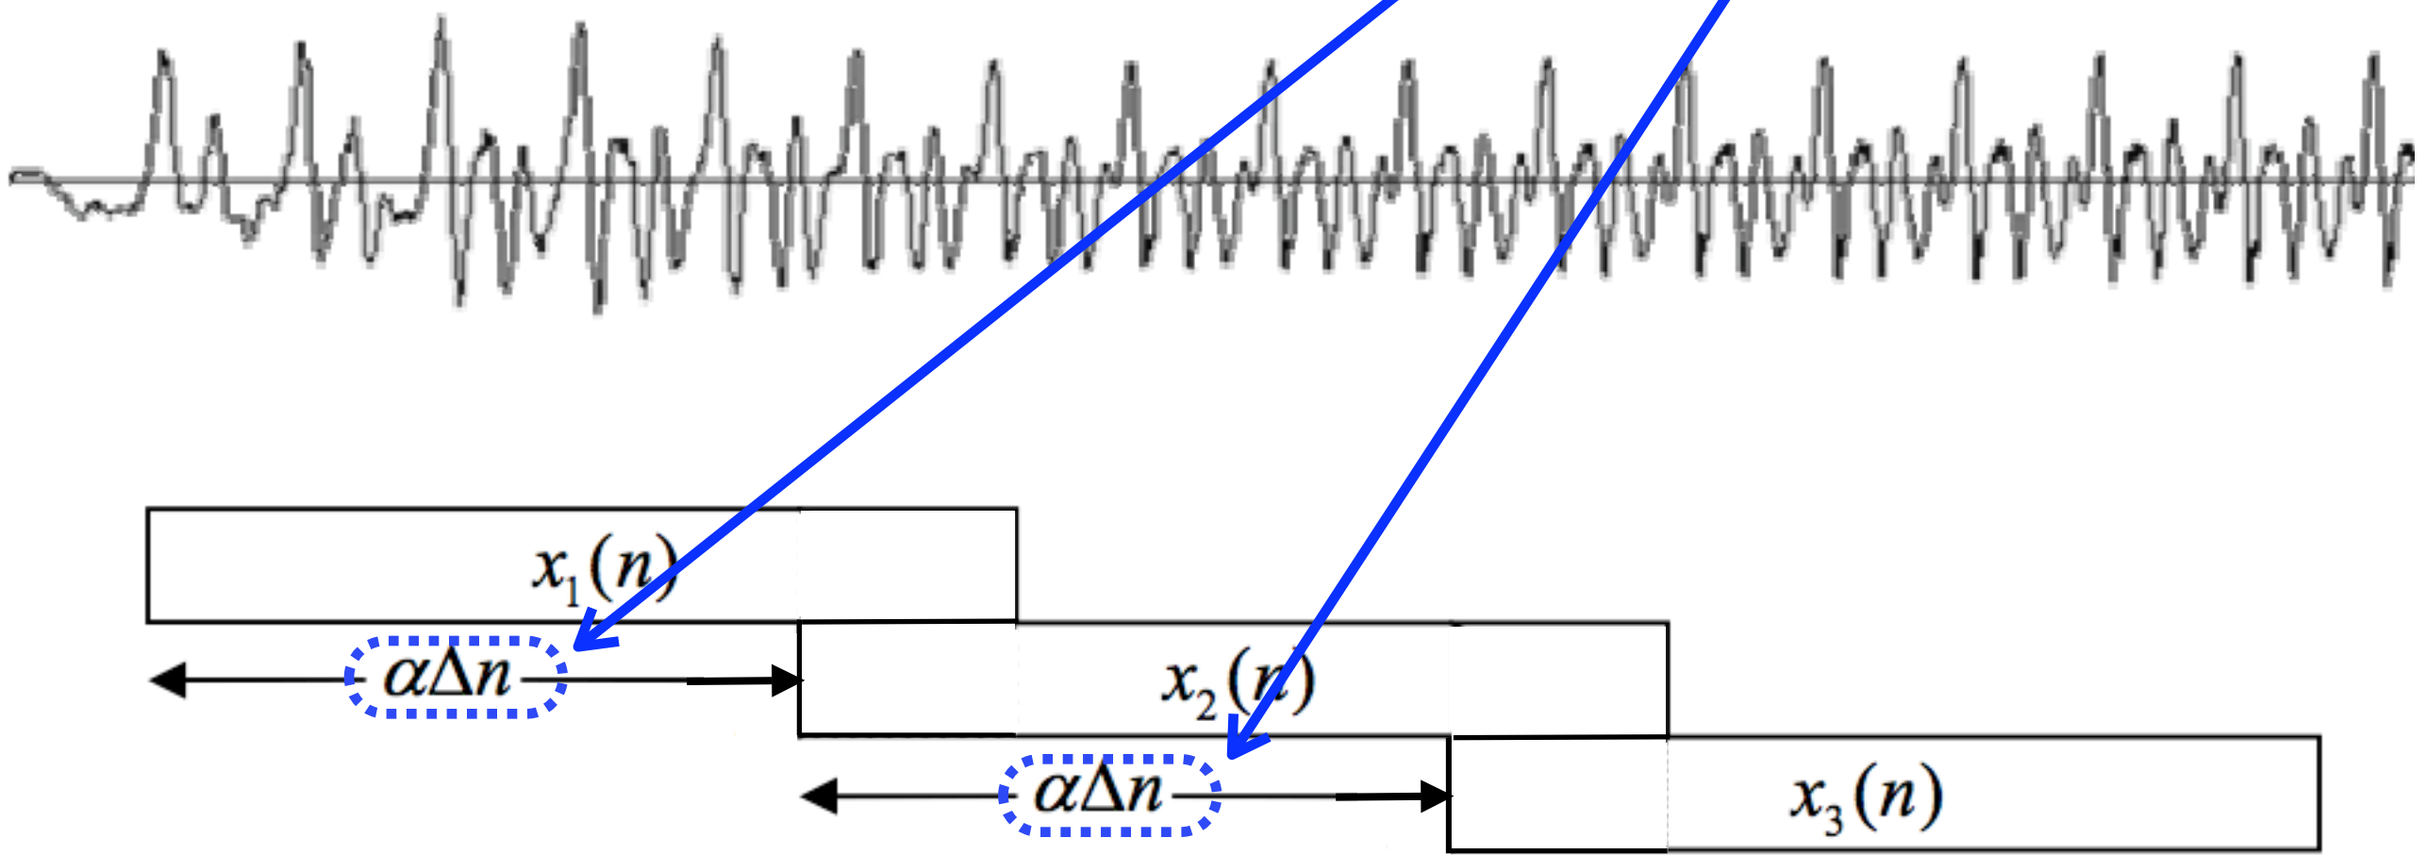

"Synchronous OLA" - SOLA

Step 3: Tune each Δt_i to minimize artifacts, and then create final waveform by summing all blocks, doing a crossfade at the overlaps.

"overlap and add" - OLA

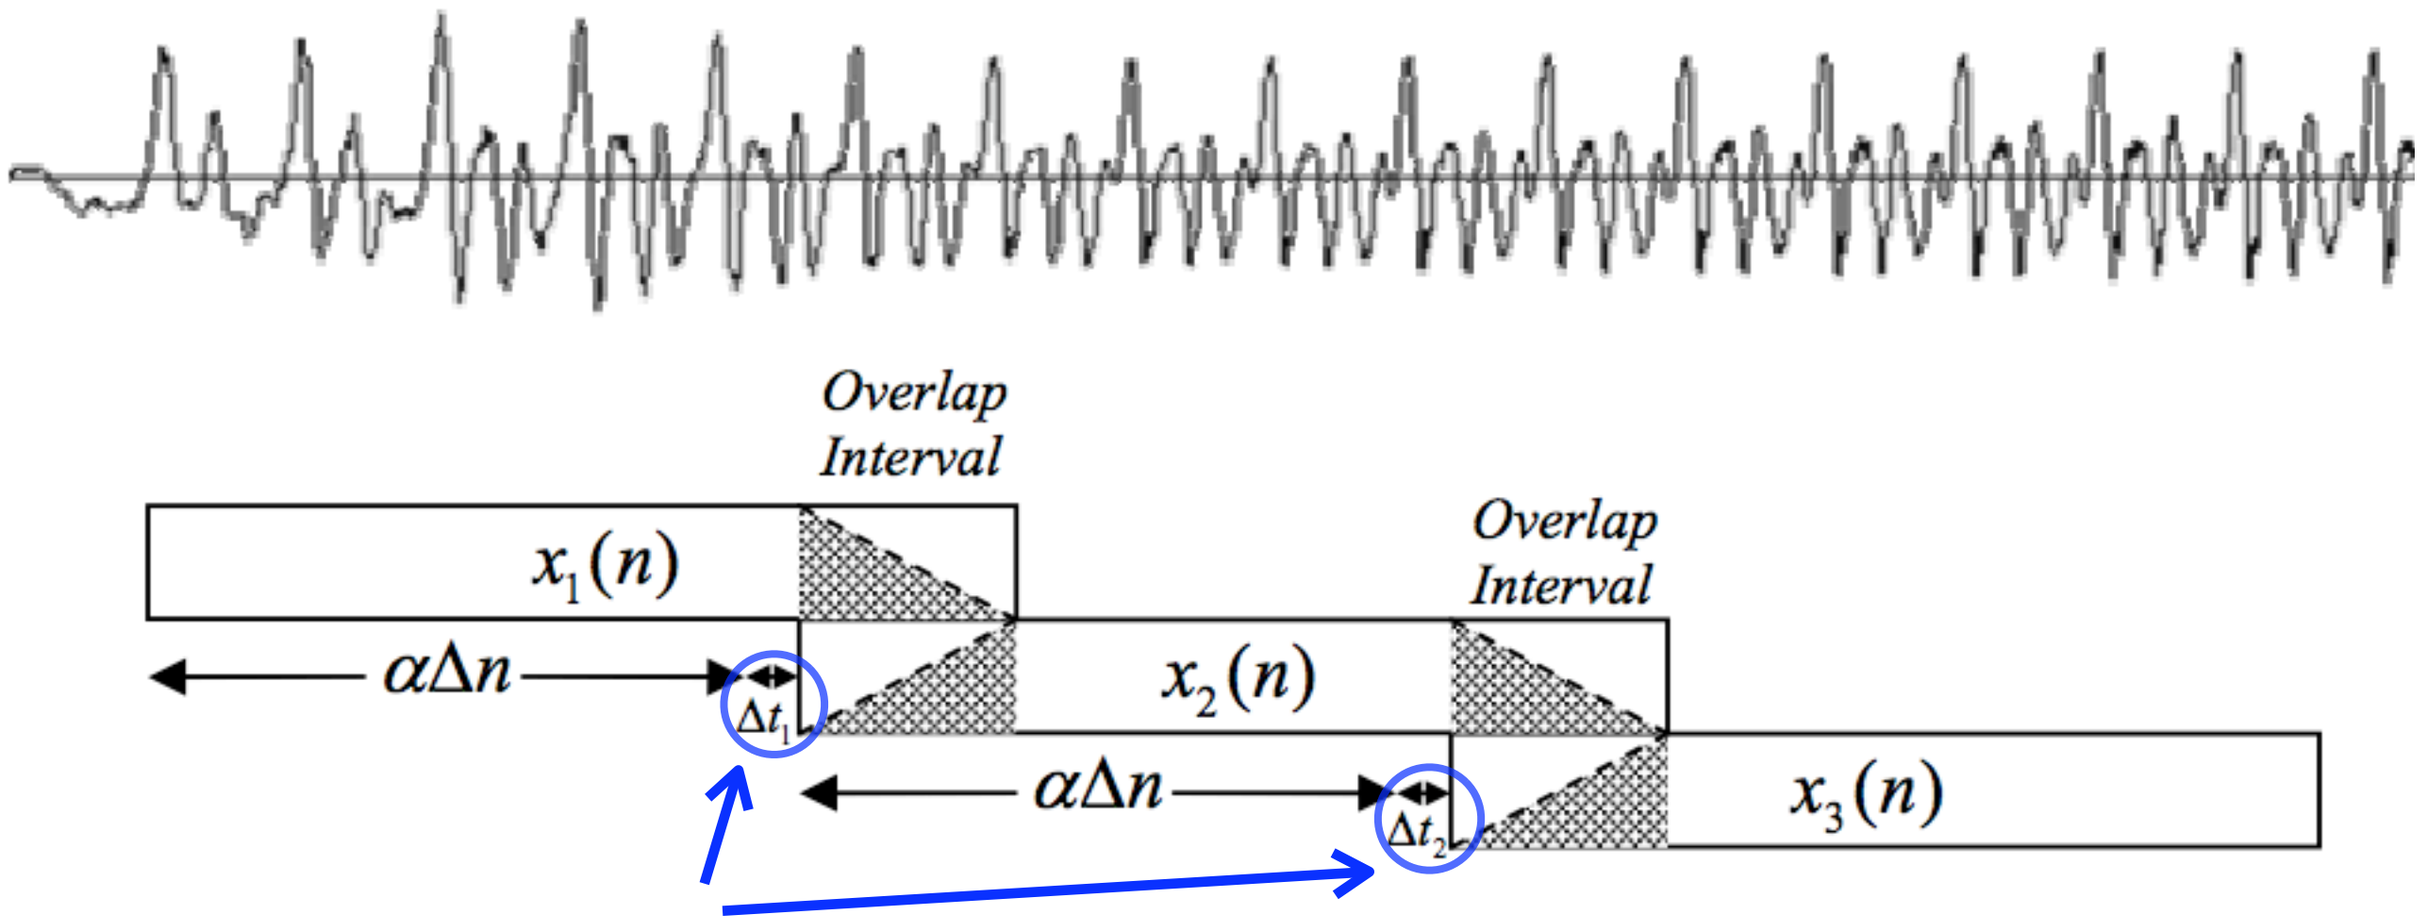

Tuning Δt_i is primary way methods differ.

PSOLA, WSOLA, PICOLA, ...

Another trick: Detect transients, don't OLA them.
First use: Lexicon 2400, 1986.

Example: PSOLA time-stretching

First, find the periods of the (speech-like) waveform.

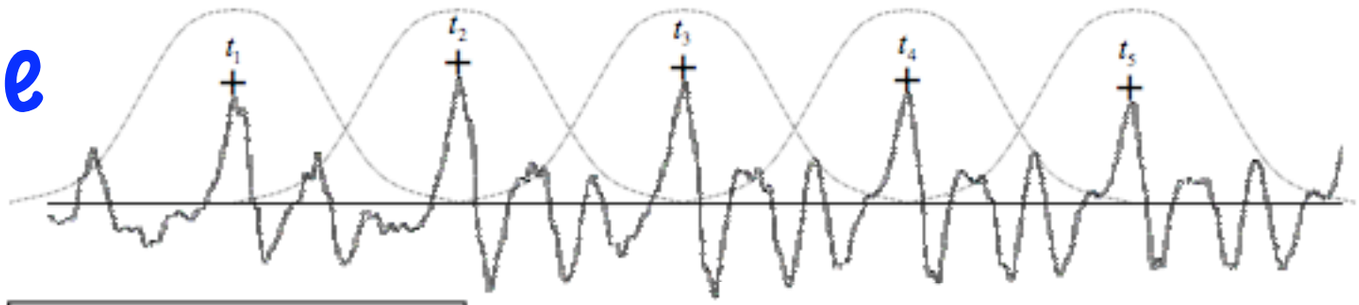

Make each block exactly two periods in length.

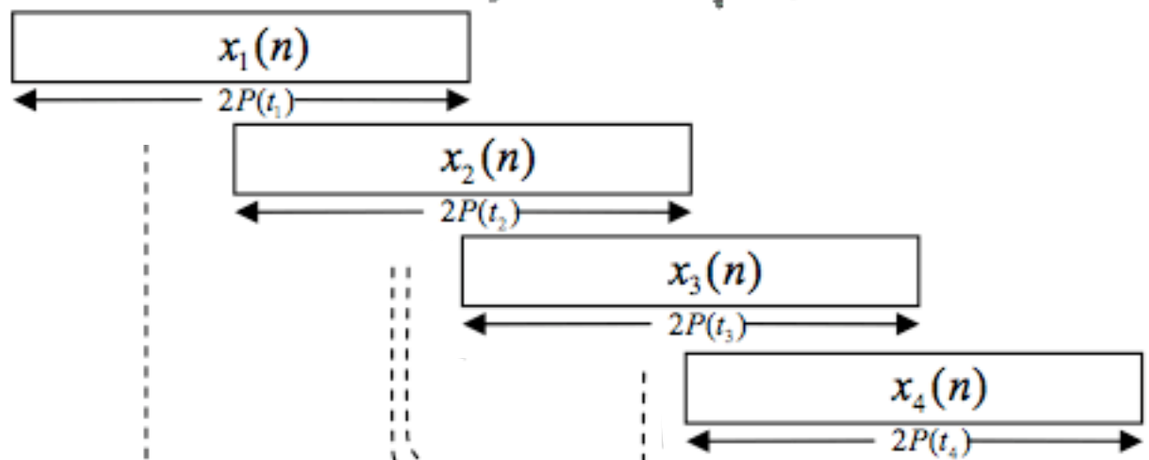

Stretch by duplicating pitch periods.

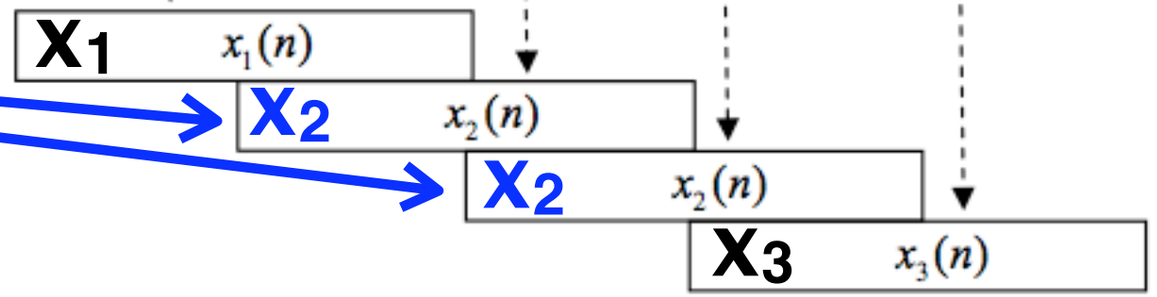

Crossfade blocks for final output

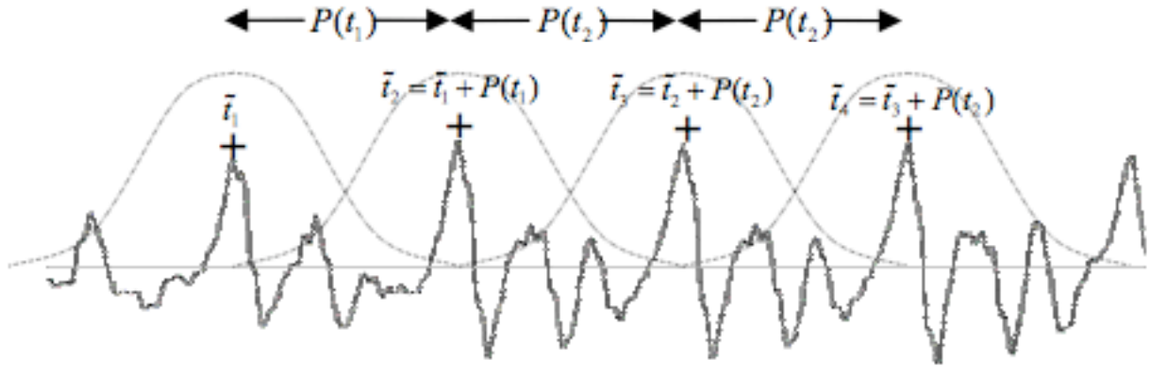

How do SOLA algorithms sound?

PICOLA: SOLA variant that uses waveform correlation to fine-tune Δt_i. Does not detect transients. As implemented in sfront.

- **Trumpet sample:** original, 0.5α, 2.0α. Play

- **Synthesized bass and nature sounds:** original, 0.5α, 2.0α. Play

- **Excerpt from Giant Steps (John Coltrane):** original, 0.95 α, 1.05 α. Play

Next Class: Spectral methods.

Project Idea

“Echos of George Bush”

The George W. Bush Public Domain Audio Archive

Brought to you by The Bots

The George W Bush Public Domain Audio Archive is a public domain database of the speeches of George W. Bush. Every phrase from each major speech has been made into an individual audio file, where the filename is, in most cases, the exact text content of the sample. This allows you to search the entire database for individual keywords. In the list of links below, "Individual files" links allow you to download a zip file containing the individual phrases. "Linear recording" links are an mp3 of the entire speech, which may be downloaded or streamed. In some cases,either the linear recording or the individual phrases are not available. The Bots have used this database to create the parody song Fuzzy Math. What you do with the database is up to you. Musicians, linguists, historians, media professionals, students and activists may all find this database to be of use. You are free to download and use these samples for any purpose, both non-commercial and commercial.

If you want to do something creative with the samples, we've been having a contest here called The Bush Raps (Re)Mix Contest. The contest is over now, but we're still adding links to new songs. Send along a link to your Bush parody mix, and we'll put it up on this site, giving you instant worldwide exposure. If you need inspiration, Fuzzy Math by The Bots is an example of what can be done with the archive. What you do is up to you. You can usually find the latest details either in bc's blog or on The Bots Forum.

The samples are arranged by individual speech in roughly chronological order. Always click the links on this page to get the files. Because of extraordinarily high downloading, the actual files will be on a rotating group of servers, so a direct link may stop working when the file is transferred to another server. The links on this page will always be updated to point to the current server.

Misc 2000 and 2001 (individual samples)

Debate With Al Gore 2000 (individual samples)

Oath of Office 20 Jan 2001 (linear recording)

Speech on 11 Sep 2001 (individual samples)
Speech on 11 Sep 2001 (linear recording)

Speech on 12 Sep 2001 (individual samples)
Speech on 12 Sep 2001 (linear recording)

Basic Idea: Voice samples echo MIDI phrases

Piano improvisation: Left-hand chords,
right hand monophonic melody line phrases.

Running MIDI analysis to detect phrase completion.

A George Bush phrase is located in the database that "echos" the piano line. It is repeated several times, with signal processing that evokes "echoing".

Choosing a voice that echos a melody

- **Voice phrase length == MIDI phrase length.**
- **Vocal cadence rhythm == MIDI phrase rhythm**
- **Vocal pitch track evocative of MIDI melody.**

Also: Verbal continuity. Ideally, the concepts spoken in each phrase should topically connect ...